ENTROPY ECONOMICS
GLOBAL INNOVATION + TECHNOLOGY RESEARCH

BRET SWANSON
PRESIDENT

February 18, 2015

Ms. Marlene H. Dortch
Secretary
Federal Communications Commission
445 12th Street, SW
Washington, D.C. 20554

Re: Protecting and Promoting the Open Internet GN Docket No. 14-28; *Framework for Broadband Services* GN Docket No. 10-127.

Dear Ms. Dortch:

Please find attached several reports that may be relevant to the Commission's Open Internet and Broadband proceedings.

The first item is "Digital Dynamism: Competition in the Internet Ecosystem," a review of the complex web of technical and commercial relationships across the Internet economy.

The second item, "Hyperconnected: The New Network Map," is a graphical representation of the monumental shifts in communications networks over time.

The third item is a report called "How the Net Works: A Brief History of Interconnection."

Sincerely,

Bret T. Swanson

Digital Dynamism

Competition in the Internet Ecosystem

November 2013

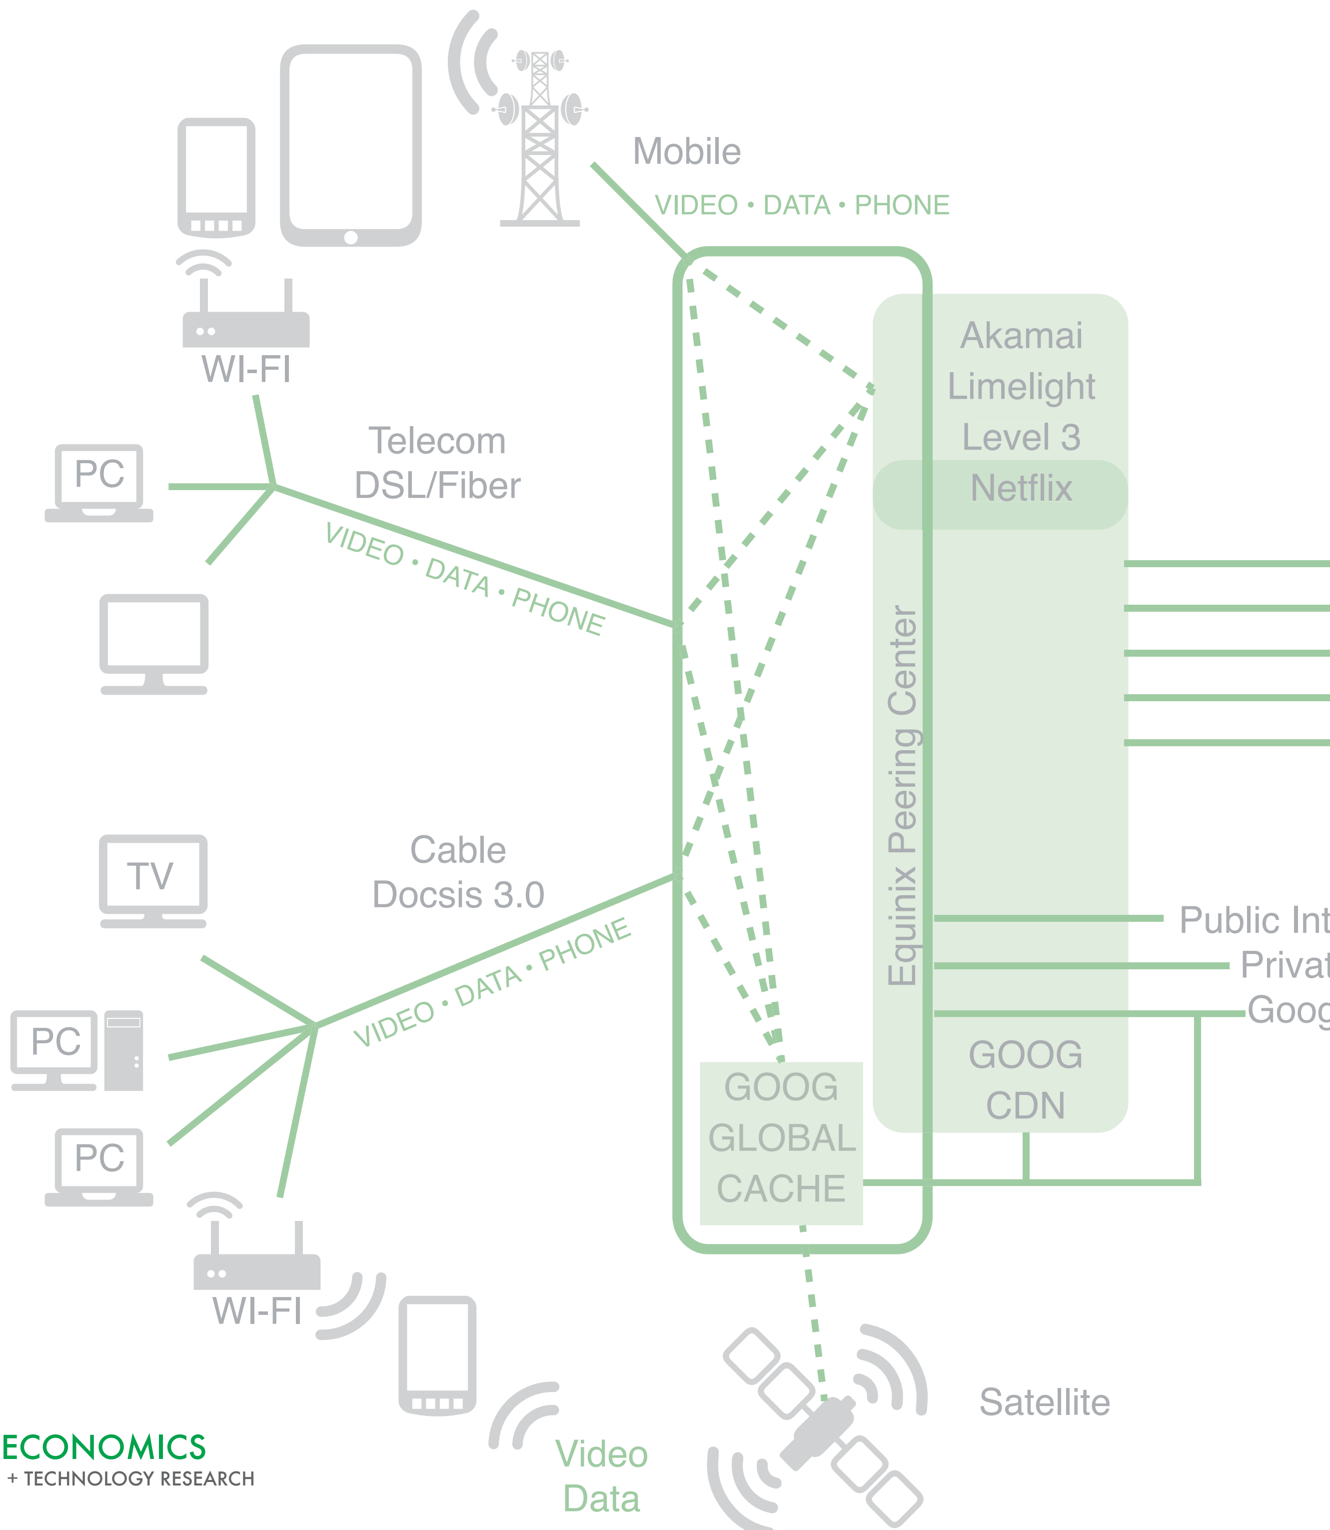

ENTROPY ECONOMICS
GLOBAL INNOVATION + TECHNOLOGY RESEARCH

Digital Dynamism
Competition in the Internet Ecosystem

November 12, 2013

The Internet is altering the communications landscape even faster than most imagined.

Data, apps, and content are delivered by a growing and diverse set of firms and platforms, interconnected in ever more complex ways. The new network, content, and service providers increasingly build their varied businesses on a common foundation – the universal Internet Protocol (IP). We thus witness an interesting phenomenon – the *divergence* of providers, platforms, services, content, and apps, and the *convergence* on IP.

The Dynamic Internet

The dynamism of the Internet ecosystem is its chief virtue. Infrastructure, services, and content are produced by an ever wider array of firms and platforms in overlapping and constantly shifting markets.

The simple, integrated telephone network, segregated entertainment networks, and early tiered Internet still exist, but have now been eclipsed by a far larger, more powerful phenomenon. A new, horizontal, hyperconnected ecosystem has emerged. It is characterized by large investments, rapid innovation, and extreme product differentiation.

- Consumers now enjoy at least five distinct, competing modes of broadband connectivity – cable modem, DSL, fiber optic, wireless broadband, and satellite – from at least five types of firms. Widespread wireless Wi-Fi nodes then extend these broadband connections.

- Firms like Google, Microsoft, Amazon, Apple, Facebook, and Netflix are now major Internet infrastructure providers in the form of massive data centers, fiber networks, content delivery systems, cloud computing clusters, ecommerce and entertainment hubs, network protocols and software, and, in Google's case, fiber optic access networks. Some also build network devices and operating systems. Each competes to be the hub – or at least *a* hub – of the consumer's digital life. So large are these new players that up to 80 percent of network traffic now bypasses the traditional public Internet backbone.

- Billions of diverse consumer and enterprise devices plug into these networks, from PCs and laptops to smartphones and tablets, from game consoles and flat panel displays to automobiles, web cams, medical devices, and untold sensors and industrial machines.

Competition and Cooperation

The communications playing field is continually shifting. Cable disrupted telecom through broadband cable modem services. Mobile is a massively successful business, yet it is cannibalizing wireline services, with further disruptions from Skype and other IP communications apps. Mobile service providers used to control the handset market, but today handsets are mobile computers that wield their own substantial power with consumers. While the old networks typically delivered a single service – voice, video, or data – today's broadband networks deliver multiple services, with the "Cloud" offering endless possibilities.

The competitive and cooperative relationships among all these companies are complex, dynamic, and multidimensional. A cable MSO, for example, which used to compete solely with broadcast TV, now competes with many more firms in many more markets.

In its traditional TV content business, satellite (Direct TV and Dish) and telecom (Verizon FiOS and AT&T U-verse) now offer the same hundreds of channels that cable offers. Telecom, 4G wireless, satellite, and even public Wi-Fi networks compete with cable in broadband data. Telecom and mobile compete with cable in services like voice, as do web players like Skype and messaging apps like WhatsApp. And the burgeoning world of web content – Netflix, Hulu, Amazon Video, MLB.com, the endless bounty of YouTube – competes with cable's traditional content.

These Internet companies, however, also offer synergistic benefits to telecom and cable firms. Netflix, YouTube, and Skype, for example, promote strong demand for broadband Internet access services. In the same way, the iPhone both challenged mobile carriers' control of the handset market and yet boosted demand for mobile services. Many firms and technologies are thus often competitors and complements at the same time.

Multidimensional Competition

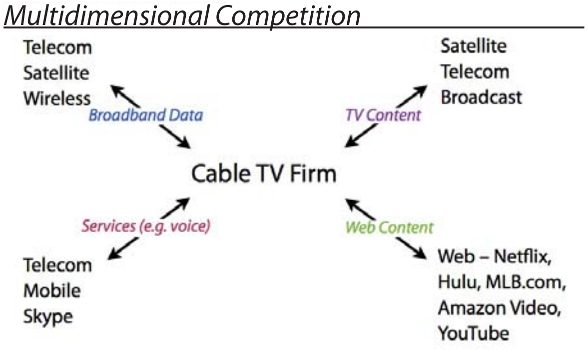

U.S. Broadband Success

The success of the U.S. broadband ecosystem suggests government policy has been mostly supportive. Light-touch or even no-touch regulation has fostered experimentation, entrepreneurship, and explosive growth in network and computer capacity and services. More than other nations, the U.S. focused on facilities-based competition. Over the past 15 years, private firms invested more than $1.2 trillion in broadband networks, and today the U.S. boasts:

- close to 90 million residential broadband subscribers, up from around five million in the year 2000;
- 327 million mobile subscriptions and 302,000 mobile cell sites, including the world's broadest deployment of 4G mobile networks and devices;
- broadband networks that are among the world's very fastest, most ubiquitous, and most robust;
- Internet and IP traffic of some 20 exabytes per month, up from just 10 terabytes per month – a two-million-fold increase in two decades; and
- the great majority of the world's most important digital innovations and firms – Google, Amazon, Salesforce.com, Twitter, mobile operating systems and millions of "apps."

Next Generation Policy

The growth, complexity, and dynamism of this market (1) expose the conterproductivity of older policies that may no longer be relevant or justified; and (2) challenge the wisdom and authority of newer attempts at top-down micromanagement of networks, digital business models, and wireless spectrum.

Today's policymakers and regulators should:

- recognize the complexity and dynamism of networks and the services that flow over them;
- appreciate the success of, and endeavor to sustain, the successful multistakeholder governance of the Internet;
- remove existing barriers to investment, and prevent the erection of new ones; and
- avoid prescriptions or proscriptions of particular business models or technical architectures that could stifle experimentation. EE

Digital Dynamism: Competition in the Internet Ecosystem

> *From Vertical Voice Networks to Horizontal Hyperconnectivity*
> *Overlapping Networks, Overlapping Businesses*
> *Cloud + Wireless = Everything Over Everything*
> *U.S. Broadband, A Success Story*
> *What Policies Will Sustain Internet Innovation?*

BRET SWANSON > November 12, 2013

The Internet is altering the communications landscape even faster than most imagined.[1] In the last two decades, U.S. Internet and IP traffic has grown to some 20 exabytes per month from just 10 terabytes per month – a two-million-fold increase. Traffic continues to grow nearly 50 percent per year.

In the last five years, the number of mobile app downloads has exploded, from essentially zero in early 2008 to a cumulative total of more than 100 billion today.

The topology of our networks is shifting, too. Data, apps, and content are delivered by a growing and diverse set of firms and platforms, interconnected in ever more complex ways. At the same time, we use the old voice network less and less every day. The new network, content, and service providers, moreover, increasingly build their varied businesses on a common foundation – the universal Internet Protocol (IP).[2] We thus witness an interesting phenomenon – the *divergence* of providers, platforms, services, content, and apps, and the *convergence* on IP.

The success of the U.S. broadband ecosystem suggests government policy has been, at least directionally, supportive. Over the last two decades, light-touch or even no-touch regulation has fostered experimentation, entrepreneurship, investment, and explosive growth in network and computer capacity and services. Yet these dramatic changes lead to new policy questions and put in stark relief older policies that may no longer be relevant or justified.

These are the chief questions of our report: What does today's Internet ecosystem look like, and how does it work? How did we get here? And what government policies are most likely to support continued investment and innovation?

The Dynamic Internet

The dynamism of the Internet ecosystem is its chief virtue. Google, Amazon, Apple, Microsoft, Facebook, and Netflix are today major Internet infrastructure companies. We used to think of them as, respectively, search, ecommerce, computer, software, social, and motion-picture-delivery firms. But today they build and operate vast data farms and fiber networks. Several build mobile devices. Several build operating systems and browsers. All are competing to be the hub – or at least *a* hub – of the consumer's digital life. Each, however, approaches the converged digital world from a different angle and with a distinct business model.

This is possible in large part because the network – the Internet – supplies a standard infrastructure that supports multifaceted content, services, and devices.

The traditional telecom companies are of course a central factor in the digital equation. Here, too, the field is shifting. Cable disrupted telecom through broadband cable modem services, but now cable is being disrupted by free content from YouTube and subscription services like Netflix. Mobile is a massively successful new business, yet it is cannibalizing wireline services, with further disruptions from Skype and other IP communications apps. Mobile service providers, moreover, used to control the handset market, but today handsets have become mobile computers that wield their own substantial power with consumers. The iPhone, in other words, reorganized the whole mobile industry. The bottom line is that the competitive and cooperative relationships among all these companies are complex and dynamic.

New Policy Temptations

The Internet arrived with force in the mid-1990s and immediately challenged the existing framework of telecommunications policy. Broadband was a new technology, a new product, and it delivered new kinds of content and services. After some initial stumbles, the U.S. got broadband policy largely right in the 2000s, and the digital universe exploded. We now enjoy fiber-to-the-home and 4G wireless, among other access technologies, all linked to the endless resources of the cloud.

With this exaflood of new technology and content, and the overall growth and influence of the digital economy, however, new sets of policy questions arise. Net neutrality, for example, seeks new constraints on network architectures and business models. The definitions of net neutrality, moreover, morph as fast as the networks they propose to regulate. Although two decades of spectrum auctions and a healthy secondary spectrum market allowed the U.S. to become the world leader in mobile innovation, wireless spectrum policy is regressing, becoming more complicated and contentious. Mandated wireless data roaming is another example of a rule beyond the framework of our old telecom laws.

It is far from clear that these new rules are wise or that authorities, such as the Federal Communications Commission (FCC), have the legal power to impose them. Even as regulators propose additional rules for the era of the broadband cloud, however, much of the old telephone regulatory infrastructure remains. We are thus layering new platforms for the regulation of the Internet on top of the largely obsolete platforms for the regulation of telephones.

Advocates of both the new rules and old rules often justify them based on a traditional view of telecom. Underlying many of these policy suggestions is a central worry – that one or two large firms might dominate the market. But does this world still exist? Just what *is* "the market"? Can any one firm "dominate" for long? And if the communications market has changed in fundamental ways, is *either* set of rules justified? In other words, might we need an even bigger, broader rethink of communications policy?

The Vertically Integrated Voice Network

The old telephone network was built to do one thing – transmit two-way voice conversations. The telephones attached at the endpoints of the network were simple, dumb devices. One company built and operated most of the network from end to end. As seen in Fig. 1, the architecture was rather simple – a vertically integrated system.

On January 1, 1984, Judge Harold Greene's order broke up this integrated system. AT&T kept the long distance network and service, while the seven new "Baby Bells" assumed control of the local networks in seven regions. These changes, however, were largely cosmetic. They did not fundamentally alter either the technology or architecture of the network or the services delivered over it. (See Fig. 2) Far more important for competition, innovation, and consumer choice and welfare would be rival technologies and non-telephone platforms, such as cable and the Internet.

Early Convergence

In the early 1990s, the new landscape began to emerge. Using dial-up modems and serv-

1 The Vertically Integrated Voice Network

The communications network is designed, built, and operated mostly by one firm. The network does one thing. The content is supplied by the end users.

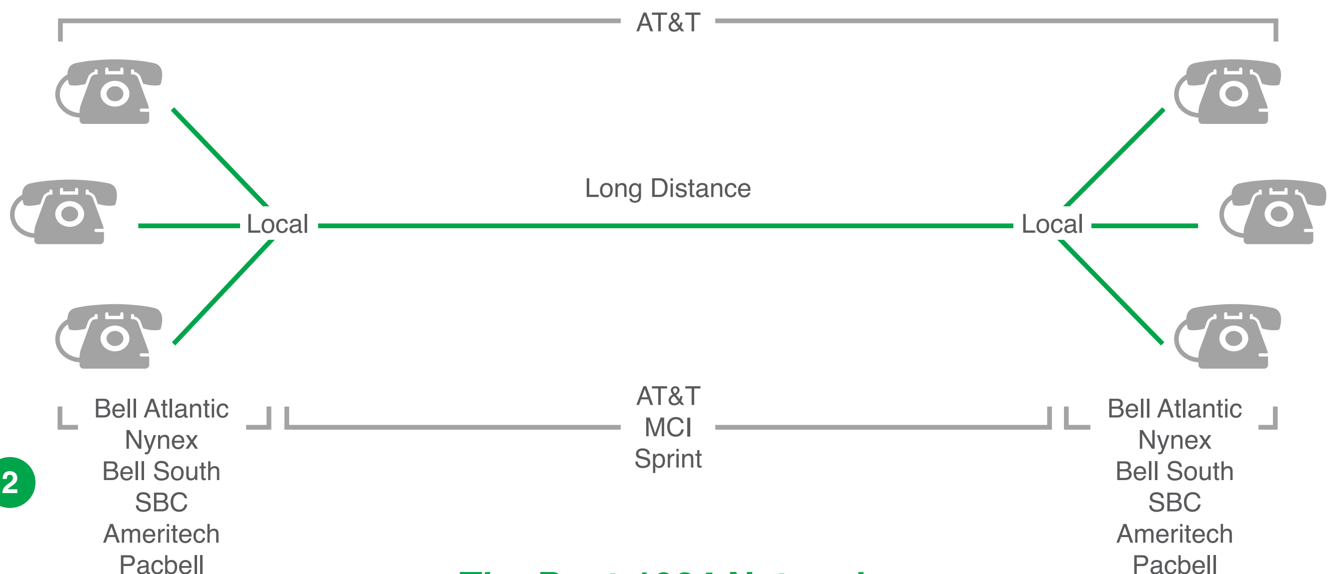

2 The Post-1984 Network

Same network, new names. The network breaks apart into long distance carriers (AT&T, MCI, Sprint) and local phone companies – the “Baby Bells.” The biggest technological innovation is fiber optics for long distance. Dial up modems offer limited access to the early Internet.

3 Early Convergence/Divergence Late 1990s – Early 2000s

More communications networks begin offering more choices. With cable modems and DSL lines, respectively, cable TV and telecom firms offer broadband access to a common, public Internet, which yields early convergence of data-based services and content. Mobile begins rapid growth.

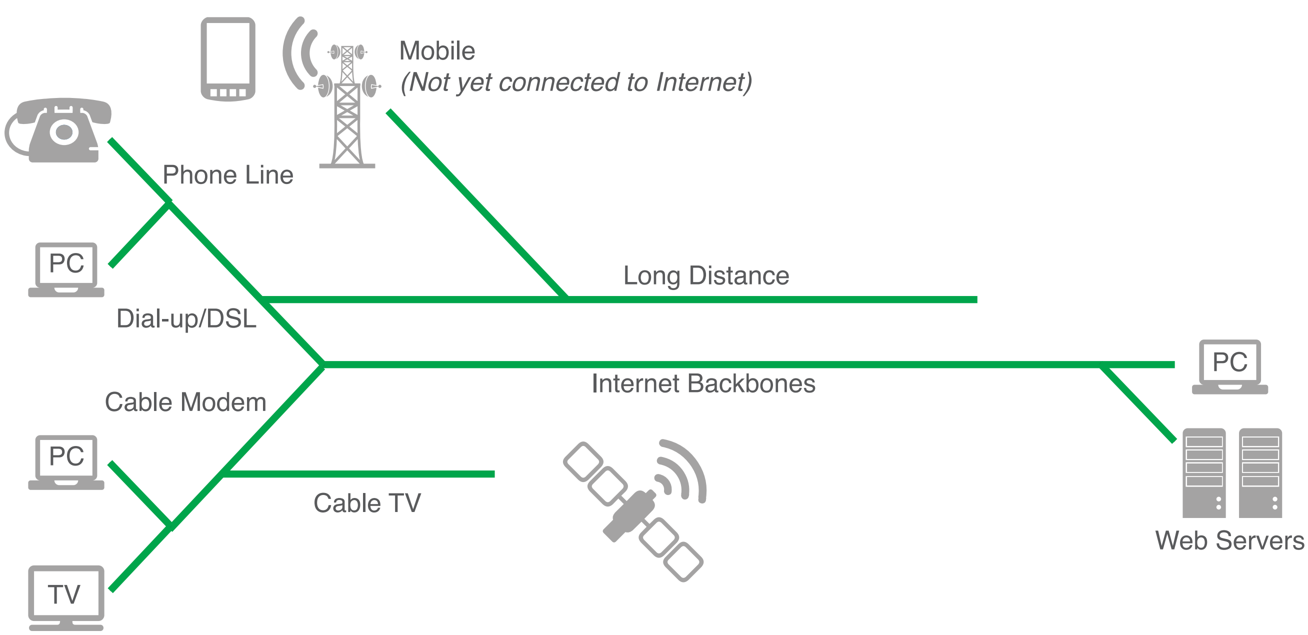

ices such as Compuserve and Prodigy, and later, AOL, a fast-growing number of average Americans with personal computers (PCs) began accessing the Internet over their phone lines.

By the late 1990s, the Web took off, dot-coms boomed, and new broadband cable modem and DSL services reached several million subscribers, though narrowband dial-up service still dwarfed the new broadband technologies. Meanwhile, firms like Level 3, Global Crossing, Williams, MCI, Sprint, Broadwing and others laid tens of millions of miles of new fiber optic cable, yielding competing Internet backbones of enormous capacity.

Here we saw the beginning of "convergence." As depicted in Fig. 3, the telecom and cable companies now offered competing Internet access services, mostly plugging into a common, public Internet backbone.

Yet in their traditional lines of business – voice and TV – telecom and cable companies still offered distinct services over distinct networks.

Mobile, satellite, and broadcast TV, meanwhile, still had almost no connection with the Internet. Mobile was still a two-way narrowband voice service. Satellite, now a competitor, instead of just a facilitator, of cable TV, was still a one-way broadcast video service.

Internet access was a new product, offering revolutionary access to information. The Internet, however, did not directly compete with voice, video, radio, satellite, or the other network services.

Exponential Digital Technologies . . .

The unrivaled, compounding power of computer and communications technologies mean today's communications networks look little like those of the past. Moore's law of computers, and its corollaries for digital storage and bandwidth, are at the heart of today's new competitive landscape.

In the middle of the century, Bell Labs – the technology arm of AT&T – invented the future of communications. Just as Claude Shannon was defining the mathematical foundations of information theory, engineers down the hall were inventing the revolutionary tools – the semiconductor transistor and the laser – that would extend Shannon's ideas into vibrant reality.

Those Bell Labs engineers assembled the original 1947 transistor by hand. Today, according to Intel, more than 100 million 22-nanometer transistors can fit on the head of a pin.[3] A new Nvidia graphics chip contains more than seven billion transistors, and semiconductor fabs worldwide now manufacture annually some one quintillion (10^{18}) of these digital switches.

Bell Labs was seeking a way to make telephone network switching more efficient. It succeeded, but it accomplished much more. Succeeding breakthroughs – combining advances in semiconductor materials, quantum electron behavior, and manufacturing miniaturization – launched the nascent computer into a whole other orbit, indeed another galaxy of possibilities. Sixty years later, the exponential computer curve of Moore's law continues.

. . . Boost Experimentation & Competition

A dozen years after the arrival of the transistor, Charles Townes, in 1960, helped invent the laser (or maser, as it was first called). Where the purity of semiconductor materials enabled digital switching and storage of information at lower power, pure laser light revolutionized the ability to transmit information over long distances – and to store and read information over long spans of time – also at low power.

At the time, copper wire carried voice conversations at a bandwidth of around three kilohertz (kHz) and could be pushed to carry 14 kilohertz signals. Today, advances in digital signal processing have, in the case of

VDSL2, expanded the bandwidth of copper wire to 30 megahertz (MHz).

Yet lasers promised potential bandwidth in the terahertz (THz) range – perhaps a million times the capacity of copper. Combined with fiber optic advances by the glass experts at Corning in the 1970s and 80s, lasers set the stage for competition in the telephone industry, most memorably MCI's and Sprint's fiber optic challenge to AT&T's long-distance service in the 1980s.

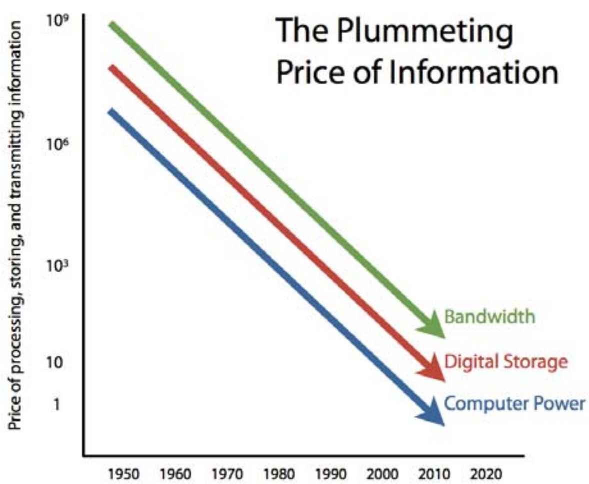

Last year, NEC and Corning unveiled an experimental fiber optic link said to transmit 1.050 petabits per second (10^{15}) over a distance of 50 kilometers.[4] That is nearly a trillion-fold leap from the old standard 3-kilobit telephone lines.

Data storage, on both spinning magnetic disks and silicon memories, has followed a similarly exponential cost-performance curve. In 1956, IBM unveiled the first commercial disk storage system, a hulking set of fifty 24-inch plates that stored 5 megabytes and sold for roughly $500,000 (in current dollars). But today one can find a 3.5-inch, 2-terabyte Seagate drive for $106.99. That is an improvement factor of some 20 million – and far more if you consider size, versatility, and reliability. Flash memory, which is more compact though not as cheap as hard disks, has revolutionized small computers, such as smartphones and tablets. And now, after many hopeful years, silicon solid state drives (SSDs) may be close to catching and overtaking hard disk storage for some applications.

When general purpose technologies like the silicon transistor and the silica-encased laser produce tools many millions of times better than the old ones, they don't just make existing infrastructure and services more efficient. They completely upset the industry landscape and, with cascading exponential ripples of new technologies, firms, and applications, create whole new industries and lift the entire economy. They are particularly open to, and supportive of, creative entrepreneurs, who can experiment with the powerful new tools and challenge existing firms, business models, political establishments, and regulatory hierarchies.

The Internet is itself a general purpose technology. It is a conceptual framework for communications constructed with these silicon and silica building blocks, with software linking them together, and content supplied by people (and now machines) across the globe.

Horizontal Hyperconnectivity

The simple, integrated telephone network, segregated entertainment networks, and early tiered Internet still exist, but have now been eclipsed by a far larger, more powerful phenomenon. A new, horizontal, hyperconnected ecosystem has emerged. It is characterized by large investments, rapid innovation, and extreme product differentiation.

- We now enjoy at least five distinct, competing modes of broadband connectivity – cable modem, DSL, fiber optic, wireless broadband, and satellite – from at least five types of firms. Widespread wireless Wi-Fi nodes then extend these broadband connections.
- Firms like Google, Microsoft, Amazon, Apple, Facebook, and Netflix are now major Internet infrastructure providers in the form of massive data centers, fiber networks,

4

Today's Hyperconnected Network

Five modes of broadband access are supplied by five types of communications service providers. Software, hardware, content, and retail companies become major Internet infrastructure providers, with massive networks and cloud computing capacity, often disrupting older services and media. Mobile devices and the App Economy achieve hypergrowth.

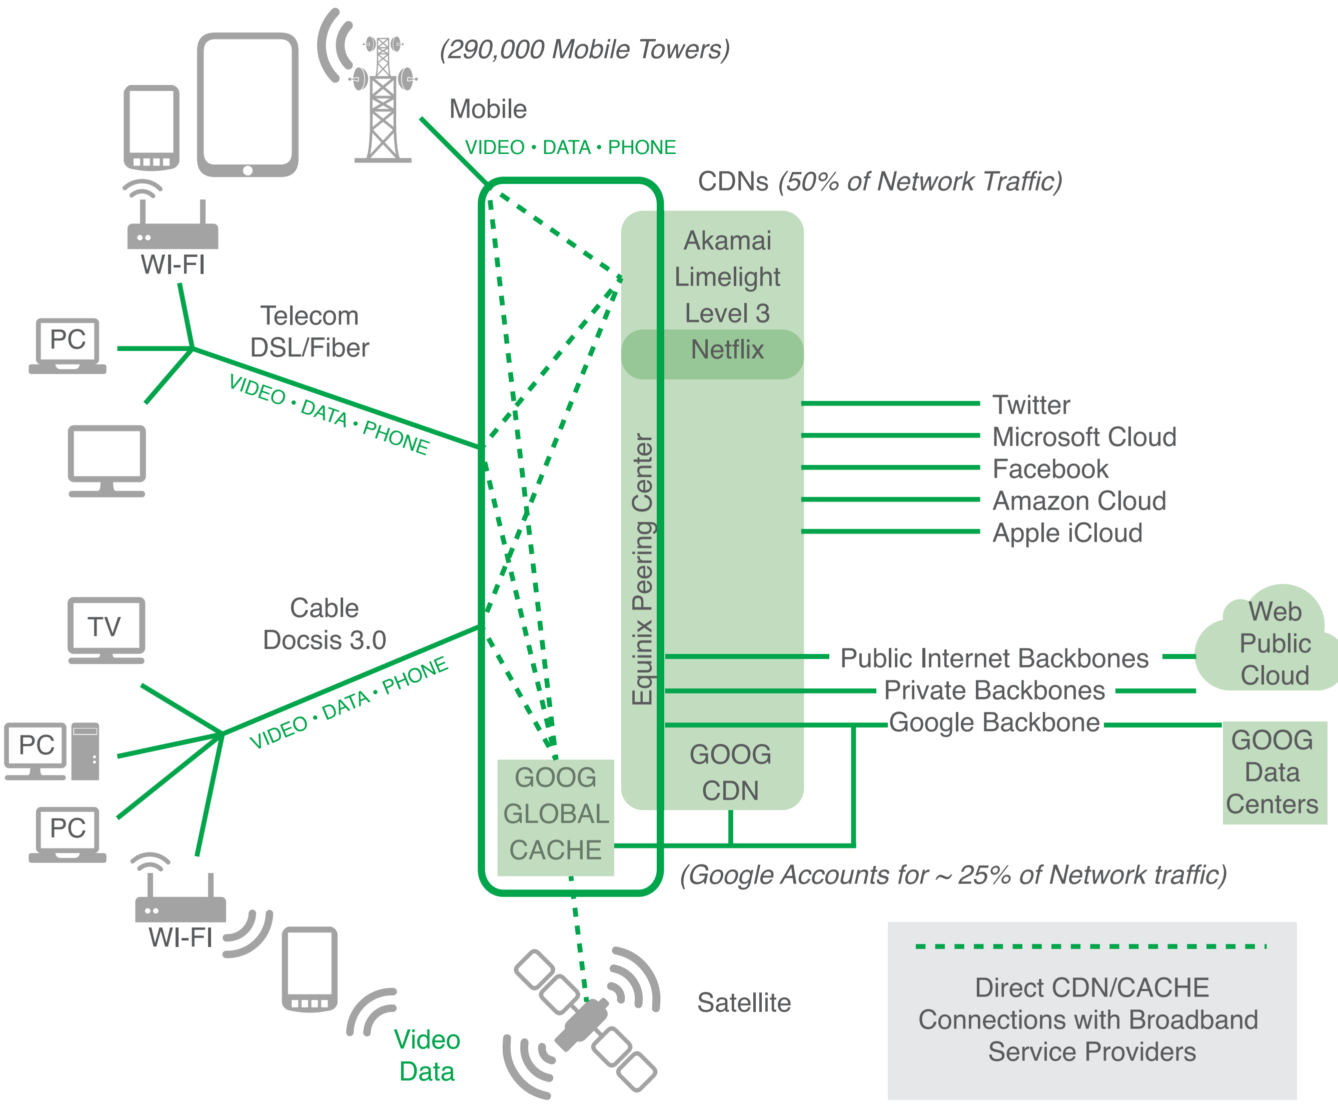

content delivery systems, cloud computing clusters, ecommerce and entertainment hubs, network protocols and software, and, in Google's case, fiber optic access networks.

- A wide range of consumer and enterprise devices plug into these networks, from PCs and laptops to smartphones and tablets, from game consoles and flat panel displays to automobiles, web cams, medical devices, and untold sensors and industrial machines.

All these networks and devices, moreover, connect in an increasingly complex web (see Fig. 4).

The topology of the Internet looks wildly different than it did just a decade ago. As Christopher Yoo, author of *The Dynamic Internet*,[5] reminds us, Internet access used to consist of a rather simple three-tier structure: access lines, regional ISPs, and backbone networks. A typical Internet session, say, sending an email or retrieving a webpage, would take the following route:

dial-up access line (tier 3) ⇆ regional ISP (tier 2) ⇆ public Internet backbone (tier 1) ⇆ regional ISP (tier 2) ⇆ DS3 access line (tier 3).

This simple formula no longer holds. Today, many networks peer directly with each other. They do so, moreover, under a variety of business arrangements, including paid peering, paid transit, and content delivery services. Comcast or Verizon, for example, may peer directly with Facebook's massive cloud infrastructure. Netflix, using its own content delivery network or similar services from Akamai, may plug in directly to AT&T's or Time Warner's broadband network. Google, naturally, plugs directly into everyone's network via its geographically distributed data farms to deliver the fastest, most reliable services (search, Gmail, maps, etc.).

Network scientist Craig Labovitz was among the first to document the growing size and power of these new Internet infrastructure players.[6] He called them "hyper giants." Indeed, by some estimates, 80 percent of today's network traffic bypasses what we used to think of as the public Internet backbone.

Early last decade, as Google's search service and advertising platform achieved global preeminence, the company realized it needed more than search algorithms and servers hosted in someone else's data farm. It needed quicker, more reliable access to end-users who wanted search answers immediately. Google's research showed that users valued quick search results more than anything. Google needed less latency, fewer hops. It needed its own global infrastructure.

So instead of operating passive servers at the end of the long ISP-backbone chain, in which data might touch a dozen or more network nodes, or hops, Google spent many billions of dollars building its own geographically distributed data centers and content delivery networks that plug in directly to the broadband service provider access networks.

Akamai had, since the late 1990s, been optimizing performance for dot-coms and content providers who wanted to reach consumers faster and more reliably. Through its network of tens of thousands of distributed "caches," it stored copies of popular pages, advertisements, banner art, and other items closer to end users.

As web video exploded in the mid-2000s, the content delivery networks (CDNs) of Akamai, Google (owner of YouTube), Limelight, and others grew accordingly. By 2010, according to network scientist Craig Labovitz, CDNs generated nearly 40 percent of all IP traffic, and today, CDNs may generate more than 50 percent of IP traffic. Today, Google alone may account for 25 percent of North American IP traffic.[7]

Like the rest of the arena, the content delivery market is highly dynamic. Soon after Netflix launched its wildly popular streaming service, it became Akamai's largest CDN customer. But just as quickly, Netflix realized it needed its own CDN infrastructure to truly optimize the user experience and has now transferred most of its video streaming to its own distributed infrastructure.

Microsoft likewise has spent billions of dollars on its own cloud infrastructure that powers its Bing search engine and its MSN, Xbox Live, Azure, and Outlook.com cloud services. Facebook had to build its own infrastructure to serve up billions of rapidly churning social network updates and to store hundreds of billions (many petabytes worth) of uploaded photographs. Apple did the same for its iTunes and iCloud services, including the App Store.

Amazon, meanwhile, leveraged its own infrastructure, which processes millions of retail transactions and user suggestions each day, to deliver a new service of outsourced cloud computer power to developers needing metered wholesale digital horsepower. Amazon also partnered with Sprint to deliver its Kindle ebooks via its Whispersync wireless service and now is challenging Netflix with its Amazon Video service.

All of these networks are of course closely coupled with the mobile infrastructure, which is ever more reliant on robust cloud services to deliver computer power and storage to "thin client" mobile devices.

These networks are linked together under a variety of technical and business arrangements.[8] Large networks may peer with one another, or exchange traffic at no cost. But networks and content providers may also use paid peering and paid transit to improve performance and more effectively access larger networks. Such peering, or interconnection of networks, often happens at the neutral hubs of Equinix, which offers data center and exchange point services in 31 markets, including 13 in the U.S. Equinix, which builds huge high-tech warehouses with access to megawatts of electricity, boasts connectivity with 900 networks, 300 cloud service providers, 500 IT service providers, and 450 financial firms. In all, according to Packet Clearing House, in 2011 there were more than 5,000 ISPs that formed "the Internet."

Every Service Over Every Network

The network is even more complex than this superficial picture. Hundreds of important players provide key hardware and software inputs that make the Web work. Yet, as depicted in Fig. 4, even the few developments highlighted here show the network is flatter, vastly more interconnected, more dynamic, more competitive, and more complicated than ever.

This generalized broadband IP network has driven – and been driven by – an increasingly generalized market for content, services, and applications. Although most of the old dedicated networks still exist, almost all forms of content and services – radio, video, voice – now also flow over the Internet. Many entirely new forms of content and services, from webpages and user-generated video to Twitter and Salesforce.com, do as well. New apps, products, sales channels, online communities, and content emerge all the time.

Usage of the new communications channels is widespread and deep. The U.S. today boasts:

- close to 90 million residential broadband subscribers, up from around five million in the year 2000;
- 327 million mobile subscriptions – or more subscriptions than people; 302,000 mobile cell sites; and the broadest 4G deployment;
- more than 34 million satellite TV subscribers, with access to more than 200 channels, plus new, better broadband Internet services;
- more than 25 million satellite radio subscribers, with access to 165 channels;

The apps, content, and communications flowing over these networks are growing fast and are increasingly diverse. For example:

- Microsoft's Skype voice-over-Internet service now accounts for one-third of all international voice traffic.
- Microsoft also has 48 million Xbox Live customers, 400 million Outlook.com users, and 250 million SkyDrive users.
- Apple's iTunes users are purchasing over 800,000 TV episodes and 350,000 movies per day.
- Apple recently announced new content partnerships with HBO and ESPN.
- Netflix has 40 million users who view more than a billion hours of movies and TV each

Multidimensional Competition and Cooperation

***Table 1** – On the Internet, technologies, products, and firms compete and cooperate in many dimensions, producing explosive innovation and consumer benefits*

Competition
Wired vs. Wireless
Cable vs. Telecom vs. 4G Wireless (broadband)
Cable vs. Telecom vs. Mobile (services; e.g. voice)
Cable vs. Telecom vs. Satellite (content; e.g. TV channels)
Mobile Firm 1 vs. Mobile Firm 2 vs. Mobile Firm 3 . . .
Mobile vs. Wi-Fi
Web vs. Apps
Voice vs. Skype
Skype vs. Facetime
Text vs. Voice
Chat vs. Skype vs. Social Network messaging
TV channels vs. Netflix vs. Web channels (MLB.com, etc.)
Smartphone 1 vs. Smartphone 2
Smartphone vs. Laptop
Tablet vs. PC
Cloud vs. PC
OS 1 vs. OS 2 vs. OS 3 vs. OS 4
OS vs. Browser
Browser 1 vs. Browser 2 vs. Browser 3 . . .
iCloud vs. Dropbox vs. Google ecosystem
Device + OS + App vs. Exacloud

Cooperation / Complementarity
Broadband + Wi-Fi
Smartphone + Mobile Network
Browser + Web Content
Mobile + Cloud = powerful thin client
Broadband + Netflix
Broadband + YouTube
Mobile OS + Apps
Wi-Fi + Tablet
4G Wireless + Remote Sensors, Cars, Medical Devices
Smartphone + Camera
Amazon Cloud + App Developers
Maps + App Developers
Satellite + Sports Content
Content Delivery Network + Content
CDN + Internet Backbone
Users + Broadband + Apps + Content + Cloud
Internet User 1 + User 2 + User 3 . . .
Kindle + Mobile
Search + Everything
Device + OS + App
Device + Browser + Exacloud

month. It is also producing wildly successful original programming, such as "House of Cards," and recently announced major content partnerships with Disney and DreamWorks.

- Google's Android mobile OS now powers more than a billion devices, with an additional 1.5 million activations each day.

- Dropbox, a provider of cloud storage and document- and app-interoperability features, has 175 million users.

- WhatsApp, a messaging service, has 250 million users. A similar app called Line has 200 million users.

- Facebook, which is increasingly a platform for messaging and rich content, has 665 million daily active users and 1.11 billion monthly active users, while its Instagram photo app has 100 million monthly active users.

These network markets are characterized by rapid innovation, complementary technologies and products, and intense direct and indirect competition. The parameters of complementarity and competition are many and varied. Firms and technologies cooperate and compete along many axes, which are constantly shifting. In Table 1 nearby, we list many (but by no means all) of the ways digital ecosystem firms, platforms, and technologies compete and cooperate.

On the ledger of competitive relationships, for example:

- Cable broadband competes with telecom broadband and 4G wireless broadband. For example, "Hundreds of thousands of Americans canceled their home Internet service last year," reports *The Wall Street Journal*, "taking advantage of the proliferation of Wi-Fi hot spots and fast new wireless networks that have made Web connections on smartphones and tablets ubiquitous."[9]

- Mobile phones compete with land-line phones and voice-over-IP services like Vonage and Skype. Indeed, 40 percent of Americans have given up any land-line phone and now use mobile exclusively.[10]

- Mobile computers, like smartphones and tablets, compete with PCs and laptops.

- Apple creates an entirely new market for apps, but Google quickly counters with Android, which in a matter of just a few years becomes an even larger mobile platform.

- Netflix competes with TV, cable TV, IPTV, pay-per-view, DVDs, and online rivals Hulu, Apple TV, and Amazon Video.

- Pandora Internet radio competes with over-the-air radio and satellite radio.

- Native apps, like Microsoft Office or iOS or Android mobile apps, compete with Web-based services and apps.

- Facebook "dominates" the social network world, with more than a billion users. Yet its very size discourages some users who desire more intimate (or more private) online communities, opening the market to new social network providers.

- Apple, likewise, toiled for years with a five-percent share of the PC market. Part of its narrow allure was simply that it wasn't Microsoft. Now, with Apple playing the leading role in the smartphone and tablet markets, it is experiencing this effect from the opposite side: consumers, in search of variety and novelty, are looking for devices that "aren't Apple."

Many firms and technologies may not compete directly with others – they may not be full substitutes for other products, in other words. But because of the way the industry works – with its many partnerships, overlapping technological capabilities, digital modularity, rapid innovation, ease of entry, and entrepreneurial energy – products and firms that may not appear to be direct competitors, do in fact offer partial substitutes or otherwise constrain the ability of other firms to dominate the market.

This competition is fueled, in part, by the wide range of complementary relationships, which include:

- broadband networks and rich content;
- 4G mobile networks and mobile OSes and apps;
- content delivery networks like Netflix and content providers like Disney; and
- the millions of connected users who converse and collaborate.

This highly competitive and cooperative environment is itself a platform for rapid innovation, massive investment, falling consumer prices, and rising consumer choice.

In the old telephone world, the single type of content (voice conversations) was supplied entirely by end-users (telephone subscribers) over a single network owned (largely) by one company. In the broadcast industries, the networks fed consumers mass-market content over specific, segregated channels. Radio may have partially competed with TV for the consumer's entertainment time budget. But the various networks were otherwise not in competition, nor cooperation, with one another.

Now, with only a little hyperbole, consumers enjoy "everything over everything." And because each player has only, at most, a few pieces of the puzzle, the puzzle pieces are changing shape, and the puzzle is getting larger at a rapid rate, all players remain constrained yet hungry.

To zoom in on the competitive forces, we can choose one type of firm, say, a cable MSO, and analyze the different parameters of competition it faces (and exerts). As depicted in Fig. 5, a cable firm, which used to compete solely with broadcast TV, now competes with many more firms in many more markets.

In its traditional TV content business, satellite (Direct TV and Dish) and telecom (Verizon FiOS and AT&T Uverse) now offer the same hundreds of channels cable offers.

***Fig. 5** – Multidimensional Competition*

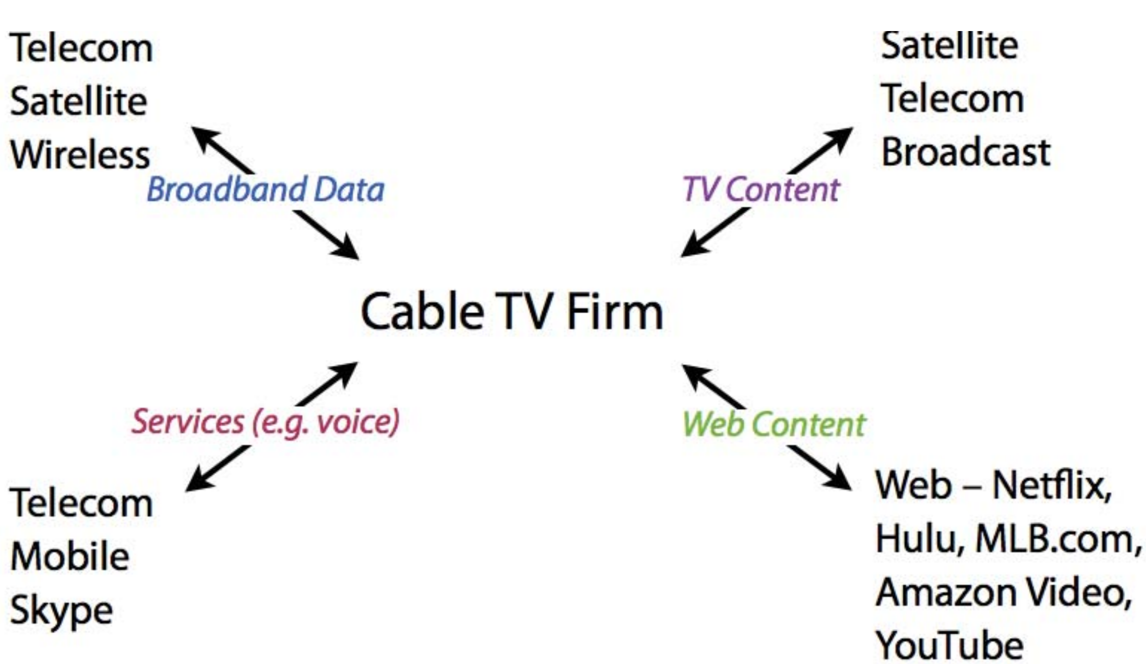

Telecom, 4G wireless, satellite, and even public Wi-Fi networks compete with cable in broadband data. Telecom and Mobile compete with cable in services like voice, as do web players like Skype. And the burgeoning world of Web content – from Netflix, Hulu, and Amazon Video to MLB.com, NBA.com, and the endless bounty of YouTube – competes with cable's traditional content.

These Internet companies also offer synergistic benefits to cable. Netflix, YouTube, and Skype, for example, promote strong demand for cable's broadband Internet access services. Many firms and technologies are thus often competitors and complements at the same time.

The Exacloud Frontier

New architectures and products will continue to challenge the ever-shifting status quo. Early this year, Otoy, a pioneering cloud graphics company, and Mozilla, the maker of the Firefox browser, unveiled a new way to bring any service, any app, to any device, regardless of platform or operating system. Using graphics supercomputers in the cloud, with petaflops of processing power, they can host any OS, app, or content and send a video stream of the "desktop view" to any device. High power 3D modeling like Autodesk's AutoCAD can thus be performed on an iPad. A Microsoft Surface or Samsung Galaxy running Android can, likewise, run Apple OS X or iOS apps. Any thin-client device can play any game, without the need to buy into a particular gaming platform or purchase a particular title. All that is required is an Internet connected device and a browser.

In July, Autodesk and Otoy unveiled their first iOS app, Autodesk Remote, which allows engineers and designers to use an iPad to access their high-powered modeling software back at the office.

Mozilla's Brendan Eich, the creator of JavaScript, calls this exacloud paradigm the future of the Web.[11] It opens a multitude of new business models for content and app providers and challenges the existing hardware-software arrangements. The exacloud's rich real-time video streams also require vast network capacity, low latency, high reliability, and closely coupled wired and wireless nodes to ensure a user experience as good as that of a client running local, native apps. Policies that encourage more investment in wired and wireless broadband are thus essential.

Policy in a Polynetwork World

The "multisidedness," modularity, network effects, and dynamic infrastructure of this ecosystem fundamentally change the way we should think about governing it.

Among those who have thought most about this new ecosystem is economist Jeffrey Eisenach, a visiting scholar at the American Enterprise Institute. In "The Theory of Broadband Competition," a detailed review of the relationships across the digital marketplace and an application of the relevant economics, Eisenach shows how the multiplicity of players and the very nature of digital technologies yield robust innovation and competition.

Eisenach summarized his findings:

> "broadband markets are now characterized, like markets in the rest of the IT sector, by dynamism, modularity, network effects, and multisidedness. The competitive dynamics of such markets are shaped by complex interactions between market-specific factors on both the demand and supply sides, but the central tendencies are straightforward.
>
> "Dynamism is shorthand for a causal circle in which firms compete by investing to create new products and, by succeeding, differentiate themselves sufficiently to earn an economic return on their investments, which attracts the capital needed to repeat the cycle. Modularity allows this process of innovation and differentiation to exploit the specialized capabilities of multiple firms to generate complementary products; it places producers of complementary goods in competition with one another over the current rents and future directions of the platforms in which they participate; and it

> creates a new type of competitor, competitive platforms, composed of loose and fluid alliances of complementers that may themselves belong to multiple platforms.
>
> "Network effects and multisidedness function in many respects as competition 'superchargers' – they magnify the effects of competitive choices through demand-side complementarities of scale and scope.
>
> "For purposes of competition analysis, broadband markets share all the key characteristics of other IT markets, including, specifically, the markets for Internet applications, content, and devices." [12]

Clearly, no rationale exists to treat these industries like the monopolies of old. The Telecom Act of 1996, and many reforms of state telecom laws since, began to acknowledge this new, more dynamic, more competitive world. Yet even the 1996 Act barely contemplated the Internet, let alone its far reaching impact.

This failure to legislate a preconceived technical or business architecture for the Internet may have been a gift. It allowed the Internet to grow up in an largely unrestrained environment, where experimentation and entrepreneurship ruled, producing explosive and mostly beneficial results. Private firms, industry groups, and scientific associations have cooperated to build the interconnection protocols and practices that make the Internet work and have served as an effective form of multistakeholder governance. The Packet Clearing House survey, for example, found that 99.51 percent of all Internet peering relationships happen "on a handshake," or customary, no-contract understandings.

Today, however, agencies are eager to regulate what have become the fastest growing networks – the broadband Internet and mobile – regardless of any explicit authority to do so.

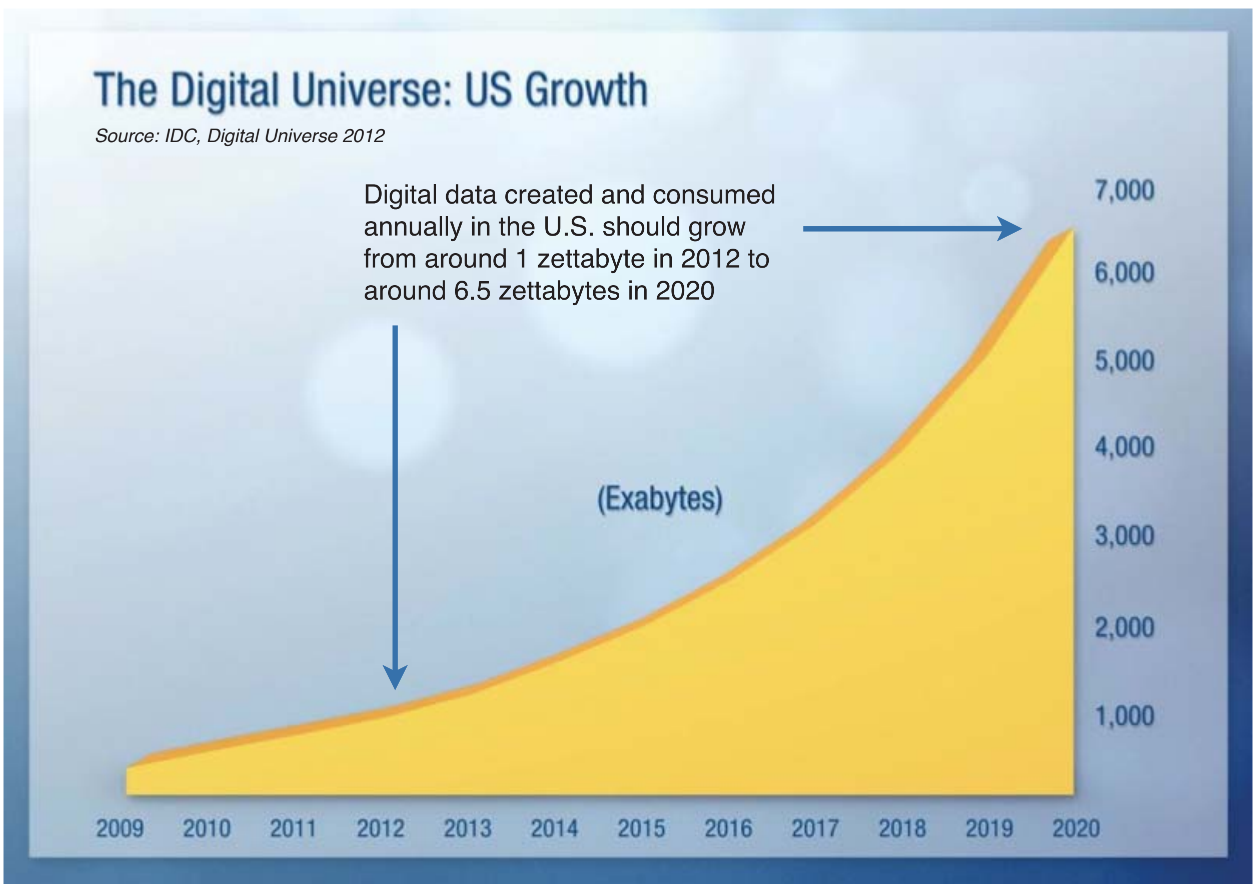

Policy today, from net neutrality to wireless spectrum, is operating in a legal realm that is simultaneously a tangled clutter and a vacuum. There is both far too much law and regulation left over from a world that no longer exists, and which is often mischievously applied to circumstances it could not contemplate; and possibly too little basic legislative and regulatory guidance for the new, highly competitive digital ecosystem.

Because the old rules are a bad fit and new rules have yet to be written, the FCC has, in a number of cases, made up new rules on its own. The Open Internet Order of 2010, also referred to as "net neutrality," regulates the technical means of managing data traffic and proscribes certain business practices and relationships. But the FCC was very creative in assuming authority over (and issuing detailed regulations governing) the Internet. Congress never gave the agency authority to do so, and firms are now challenging net neutrality in court.

Substantively, net neutrality is an example of a rule that is too pre- and proscriptive for a dynamic environment like the Internet. Had net neutrality been in force 15 years ago, important technical and business advances like content delivery networks (which deliver better network performance for a price) may have been banned and thus never developed.

The FCC's notion of net neutrality, moreover, applies the regulations to just one component of the network – broadband access. It is based on an old view of the network as a top-down, vertically integrated monopoly delivering one type of service. For instance, Google, which now accounts for 25 percent of network traffic, is largely exempt from the rules. The decidedly non-neutral, selective targeting of particular firms and network components denies the reality of a sprawling, diverse, hyper-connected system (as pictured in Fig. 4).

Wireless spectrum policy suffers from similarly old and constrained views of the ecosystem. Decades ago, the government issued spectrum licenses for TV and radio airwaves. But in an era of convergence, why should spectrum be labeled and confined to a particular technology or service? Everyone realizes this, which is why the FCC is planning incentive auctions of broadcast TV spectrum, likely to be acquired by mobile service providers.

Yet the monopoly view has infected the auction planning process. The Department of Justice and others are urging the FCC to limit the firms who can bid for and acquire these 600 MHz airwaves. In effect, the government would choose who gets the spectrum. These proposed interventions come on the heels of previous government vetoes of attempted transactions in the secondary spectrum market (AT&T's blocked purchase of T-Mobile, for example).

A third example highlights the point that policy is behind the curve. The FCC has set the goals of expanding broadband access and adoption and of transitioning from the old, limited telephone infrastructure to modern, converged, broadband IP networks. Yet a set of our broadband investors are also required by law to keep investing in the old telephone networks that the companies, and the FCC, wish to phase out.

The companies believe much of this investment is duplicative and wasteful and that it diverts capital from modern broadband. The fact is, however, that consumers and rival firms and technologies are phasing out the old telephone networks whether anyone else wishes it to happen: use of the old telephone networks is in precipitous decline. The question is whether laws and rules should deny this reality and mandate good money after bad.

Despite these wholesale changes, the old rules treat the incumbent telecom firms as if they are still monopoly providers. In many markets and for many services, however, these companies are no longer even dominant, let alone monopolistic. Wireless services are replacing many wireline offerings altogether. All-IP online offerings, such as Skype and Netflix, moreover, show how access infrastructure and access service are now often decoupled from application. According to a February 13, 2013, Telegeography report, Skype now accounts for one-third of all international phone traffic.[13]

Netflix, meanwhile, may account for one-third of U.S. broadband access network traffic during peak evening hours.[14] A new analysis now shows by the end of 2013, just one-quarter of U.S households will have a landline phone connection from an incumbent provider.[15] Yet the regulators still label these firms "dominant." Rules that presume the incumbents monopolize any component of the ecosystem – network infrastructure, access service, or applications – are outdated and have become severely counterproductive.

Incumbent providers have stated that as much as half of their wireline investments are steered into the old, increasingly obsolete networks purely for regulatory reasons. The old rules thus require that tens of billions of dollars a year be spent on infrastructure we want to retire, and that we not spend it on the networks of the future.

The government is basing many of its policies on a pre-Internet understanding of the digital ecosystem. Regulators often presume a firm's large share of a narrowly defined "market" will necessarily lead to anticompetitive behavior.

But is the government defining the market correctly? Is it accounting for the new hyperconnected ecosystem? Is it acknowledging innovation's capacity to challenge each player at every turn? Is it overestimating its ability to shape industries "better" than a natural process of innovation and competition? Is it ignoring the manifest growth and vibrancy of the industries in which it has applied the "lightest touch"? Does it consider the potentially large downside of regulation that locks in old technologies and businesses and blocks new ones? Most importantly, does it base its policies on the real world effects on consumers and the economy? Or is the government picking winners and losers? Is it regulating, in effect, not to promote basic standards and broad-based competition but to favor particular competitors and disfavor others?

The complexity and rapid innovation of the ecosystem suggest top-down micromanagement of the industry is a more difficult task than ever. The hyperconnected nature of the value chain also suggests that a policy targeting one part of the network could easily produce unintended, harmful ripple effects elsewhere, disrupting price signals and relationships. Instead of tasking centralized bureaucracies to manage specific technologies and business models, many scholars suggest we adopt a simple standard of consumer welfare.

U.S. Broadband = U.S. Innovation

In a sense, advocates of more robust centralized bureaucratic regulation of the digital economy recognize the importance of the consumer. In their case for a heavier hand, they argue that American broadband is a failure – that it is too slow, too expensive, and not widely used, especially compared to the rest of the world.[16]

Regulators should intervene more aggressively, they argue, to assist certain marketplace rivals and constrain others, hoping this will boost speeds and usage and lower costs. For a time, these arguments achieved a sort of conventional wisdom. But is this view of a sluggish American broadband economy based in fact? And would the desired policies have the intended effects – or perhaps just the opposite?

Our own analysis suggests the $1.2 trillion invested by broadband firms over the last 15 years has in fact produced networks that are among the world's very fastest, most robust, most widespread, and most used.

U.S. Broadband Speeds Among World's Very Fastest

***Tables 2-5** – Akamai's extensive global infrastructure measures actual connection speeds in real time. Its "State of the Internet" report, using four measures of access network capacity, shows some two-thirds of the world's fastest broadband networks are found in U.S. states (highlighted in green).*

Table 2 – Average Measured Connection Speed	Average Megabits Per Second
Global	3.1
1. South Korea	14.2
2. Vermont	**12.7**
3. New Hampshire	**12.0**
4. Delaware	**11.9**
5. Japan	11.7
6. District of Columbia	**11.3**
7. Utah	**11.0**
8. Hong Kong	10.9
9. Massachusetts	**10.7**
9. Virginia	**10.7**
11. Maryland	**10.6**
12. New Jersey	**10.5**
13. Connecticut	**10.4**

Table 4 – Fast Broadband Connectivity	% Above 10 Megabits Per Second
Global	13%
1. South Korea	50%
2. New Hampshire	**48%**
3. New Jersey	**45%**
4. Japan	43%
4. Vermont	**43%**
6. District of Columbia	**41%**
6. Delaware	**41%**
6. Massachusetts	**41%**
9. Rhode Island	**40%**
9. Maryland	**40%**
11. New York	**35%**
11. Connecticut	**35%**

Table 3 – Average Peak Connection Speed	Average Peak Megabits Per Second
Global	18.4
1. Hong Kong	63.6
2. Japan	50.0
3. Romania	47.9
4. District of Columbia	**47.2**
5. Vermont	**47.1**
6. New Jersey	**45.7**
7. South Korea	44.8
8. New Hampshire	**44.4**
9. Latvia	44.2
10. Massachusetts	**43.8**
10. Maryland	**43.8**
12. New York	**43.1**
12. Virginia	**43.1**
14. Delaware	**42.8**
15. Utah	**41.9**

Table 5 – Broadband Connectivity	% Above 4 Megabits Per Second
Global	46%
1. Delaware	**90%**
1. New Hampshire	**90%**
3. Switzerland	88%
4. South Korea	87%
4. Rhode Island	**87%**
6. Vermont	**86%**
7. New Jersey	**84%**
7. Netherlands	84%
9. Maryland	**82%**
9. New York	**82%**
11. Connecticut	**81%**
11. Czech Republic	81%
13. Japan	79%
14. Hong Kong	78%
14. Austria	78%
16. Massachusetts	**77%**
16. Canada	77%
16. Denmark	77%
19. Belgium	76%
19. Florida	**76%**

Source: Akamai, State of the Internet, 1Q 2013

For example, real-time speed data collected by the Internet infrastructure firm Akamai shows U.S. broadband is the fastest of any large nation, and trails only a few tiny, densely populated countries.[17] Akamai lists the top 10 nations in categories such as average connection speed; average peak speed; percent of connections with "fast" broadband; and percent of connections with broadband. Nearly all the the nations on these lists, with the exception of the U.S., are small, densely populated countries where it is far easier and more economical to build high-speed networks.

Akamai also, however, lists the top 10 American states in these categories. Because states are smaller, like the small nations that top the global list, they are a more appropriate basis for comparison. If we combine the two lists – shown in Tables 2 through 5 – we see that U.S. states dominate the overall compilation.

- Ten of the top 13 entities for "average connection speed" are U.S. states.
- Ten of the top 15 in "average peak connection speed" are U.S. states.
- Ten of the top 12 in "percent of connections above 10 megabits per second" are U.S. states.
- Ten of the top 20 in "percent of connections above 4 megabits per second" are U.S. states.

U.S. states thus account for 40 of the top 60 slots – or two-thirds – in these measures of actual global broadband speeds.

Other measures of actual network usage support these findings. For years the U.S. has generated some 60 percent more network traffic per capita and per user than Western Europe, the most comparable sample in terms of size, population, and income.[18] The newest estimates show the U.S. widening this gap, generating more than twice the per capita IP traffic of Western Europe.[19]

In a recent study of the question called "The Whole Picture: Where America's Broadband Networks Really Stand,"[20] the Information Technology and Innovation Foundation found that:

- the U.S. has the third highest rate of "intermodal competition" – access to both cable and DSL – in the OECD. Only tiny Belgium and the Netherlands enjoy more access to both cable and DSL;
- the U.S. is deploying more optical fiber than all of Europe;
- entry-level broadband prices in the U.S. are second lowest in the OECD; and
- America leads the world in 4G/LTE mobile broadband.

Other measures of the U.S. digital economy point to a healthy broadband ecosystem. The U.S. has achieved a world leading share of innovation in content, application, and infrastructure advances. From YouTube and Netflix to cloud computing, operating systems, smartphones, and tablets, the U.S. leads the way.

Broader measures support this notion as well. Economist Michael Mandel, for example, estimates the App Economy, on the fifth anniversary of the App Store, has created 752,000 U.S. jobs – up from zero.[21] All these innovations – and jobs – depend upon fast, robust broadband networks.

A New Policy Path

Given the mostly successful record of broadband access and the accompanying health of the ecosystem, re-regulation or reorganization of American broadband firms thus appears to be unwarranted.

Remember the "monopoly" threats posed by AOL and Microsoft? Or IBM? Or for that matter, the telephone company? In the late-1990s many critics even said Blockbuster was a looming monopoly in video distribution. Most often, these threats of "dominance" are not solved by intrusive policy; they are usually transcended by disruptive technology and entrepreneurial firms, even new industries. Policies that seek to constrain or reward particular firms or technologies often have the perverse impact of cementing in place the incumbent firms, technologies, and policies, far longer than is healthy.

Looking ahead, policymakers and regulators should recognize which practices have fueled broadband success, and which may be obstacles to even greater achievement. In our view, policymakers should:

- recognize the complexity and dynamism of networks and the services that flow over them;
- practice humility and restraint;
- acknowledge the multiplicity of competitive and cooperative relationships across the industry;
- remove existing barriers to investment, and prevent the erection of new ones;
- recognize the success of, and endeavor to sustain, the successful multistakeholder governance of the Internet;
- avoid prescriptions or proscriptions of particular business models or technical architectures that could stifle experimentation and investment; and
- instead look to a standard of consumer welfare, which looks at whether firm practices impose particular harms.

Jonathan Nuechterlein and Phil Weiser, authors of the authoritative communications policy book *Digital Crossroads*,[22] summed up the attitude policymakers should adopt as one of humility. "With every important decision," they advise, law makers and regulators should

> "remember[] the many times in which other policymakers have been flatly wrong in their predictions of how the telecommunications market would take shape and in their assessments of the regulatory measures needed to enhance consumer welfare within that evolving market. Humility also reminds policymakers that, over the long term, the unintended, undesired consequences of regulation can dwarf the intended, desired outcomes. That fact is not a reason for doing nothing when action is needed to correct genuine market failures. But it is a reason for policymakers to respect the market's ability to enhance consumer welfare and, as they evaluate the predicted benefits of their own regulatory involvement, to give due regard to the unpredictable course of technological and economic change."

The best approach for digital governance, as with most arenas of policy, is likely to be "simple rules for a complex world."[23] EE

---

[1] The author would like to acknowledge and thank Broadband for America for supporting the research in this report.

[2] For our purposes, "IP" is a generic reference. It encompasses a range of modern, interoperable technologies and protocols for transmission of information over data networks. It is not necessarily a specific reference to the TCP/IP protocol, although TCP/IP makes up the greatest portion of "IP." Consistent with our view of a highly dynamic Internet, new technologies and protocols will come along. See, for example, the existing UDP protocol and Google's new experimental QUIC protocol. The "IP" reference is not intended to cement in place any particular protocol or technology.

[3] Intel. "Fun Facts: Exactly How Small (and Cool) is 22 nanometers?" http://www.intel.com/content/www/us/en/silicon-innovations/standards-22-nanometers-fun-facts.html

[4] "NEC and Corning Achieve Petabit Optical Transmission." Optics.com. January 22, 2013. http://optics.org/news/4/1/29

[5] Yoo, Christopher S. *The Dynamic Internet*. The AEI Press. Washington, D.C. 2012.

[6] Labovitz, Craig, et al. “2009 Annual Report.” Atlas Internet Observatory. Nanog 47. 2009. http://www.nanog.org/meetings/nanog47/presentations/Monday/Labovitz_ObserveReport_N47_Mon.pdf; also see Craig Labovitz. “Internet Traffic Evolution: 2007-2011.” Global Peering Forum. April 6, 2011. http://www.monkey.org/~labovit/papers/gpf_2011.pdf

[7] McMillan, Robert. “Google Serves 25 Percent of North American Internet Traffic.” Wired.com. July 22, 2013. http://www.wired.com/wiredenterprise/2013/07/google-internet-traffic/ ; and Craig Labovitz. “Google Sets New Internet Record.” http://www.deepfield.net/2013/07/google-sets-new-internet-record/

[8] For a comprehensive treatment of peering and interconnection, see William B. Norton’s guide “The Internet Peering Playbook, 2013 Edition”; and his website, drpeering.net. Also useful on these topics are the research of Deepfield CEO Craig Labovitz (mentioned elsewhere in this report) and the reports periodically issued by Sandvine.

[9] Troianovski, Anton. “Cord-Cutters Lop Off Internet Service More Than TV.” *The Wall Street Journal*. May 29, 2013. http://online.wsj.com/article/SB10001424127887324682204578513262440196772.html

[10] Brogan, Patrick. US Telecom. https://www.ustelecom.org/blog/americans-continued-drop-landlines-2012

[11] Eich, Brendan. “Today I Saw The Future.” May 3, 2013. https://brendaneich.com/2013/05/today-i-saw-the-future/

[12] Eisenach, Jeffrey. “Broadband Competition in the Internet Ecosystem.” American Enterprise Institute. October 2012.

[13] “The Bell Tolls for Telcos?” Telegeography. February 13, 2013. http://www.telegeography.com/press/press-releases/2013/02/13/the-bell-tolls-for-telcos/index.html

[14] “Global Internet Phenomena.” Sandvine. November 7, 2012. http://www.sandvine.com/news/pr_detail.asp?ID=394

[15] Brogan, Patrick. US Telecom. http://www.ustelecom.org/blog/consumers-still-shedding-phone-lines-rapidly

[16] See, for example, Susan Crawford’s book *Captive Audience*. Yale University Press. 2012.

[17] “State of the Internet.” Akamai. 1Q 2013.

[18] See, for example, Bret Swanson. “World Broadband Update.” June 2011. http://www.bretswanson.com/index.php/2011/06/world-broadband-update/; also http://www.bretswanson.com/index.php/2012/06/the-real-deal-on-u-s-broadband/ and http://www.bretswanson.com/index.php/2010/10/international-broadband-comparison-continued/

[19] See Bret Swanson. “U.S. Share of Internet Traffic Grows.” TechPolicyDaily.com. October 10, 2013. http://www.techpolicydaily.com/internet/us-share-internet-traffic-grows/

[20] Bennett, Richard, Luke A. Stewart, and Robert D. Atkinson. “The Whole Picture: Where America’s Broadband Networks Really Stand.” ITIF. http://www.itif.org/publications/whole-picture-where-america-s-broadband-networks-really-stand

[21] Mandel, Michael. http://www.progressivepolicy.org/2013/07/752000-app-economy-jobs-on-the-5th-anniversary-of-the-app-store/

[22] Nuechterlein, Jonathan E., and Philip J. Weiser. *Digital Crossroads: Telecommunications Law and Policy In the Internet Age*. MIT Press. 2013.

[23] This phrase comes from Richard Epstein’s book *Simple Rules For A Complex World*. I elaborate on this concept in a paper, “Keep It Simple: A New Take On Regulation.” Bret Swanson. U.S. Chamber of Commerce Foundation. Spring 2013.

01

VERTICALLY INTEGRATED VOICE NETWORK

The communications network is designed, built, and operated mostly by one firm. The network does one thing. The content is supplied by the end users.

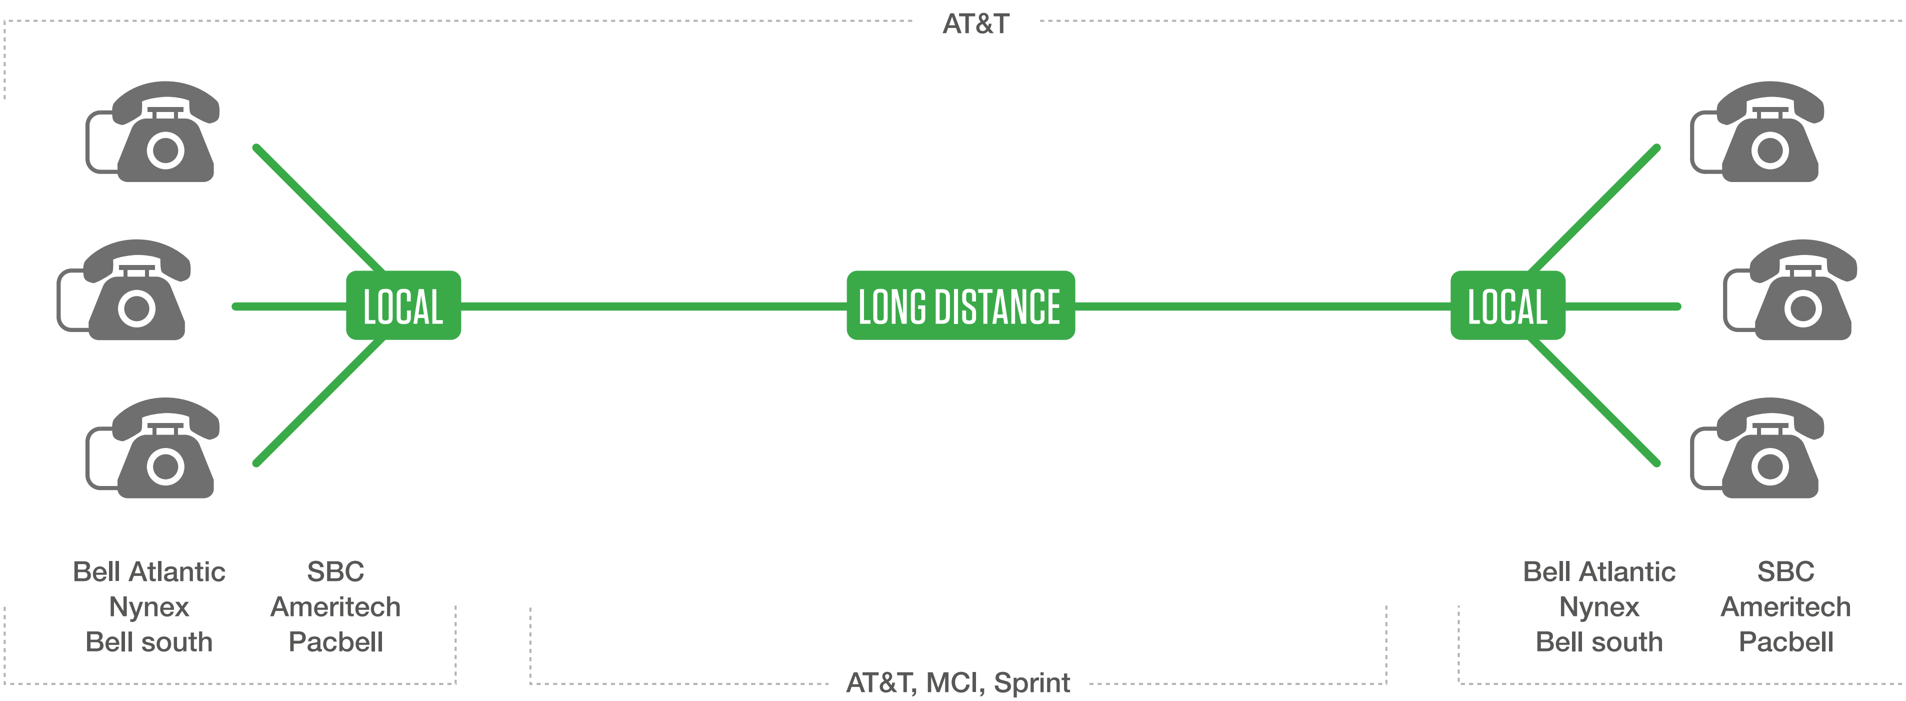

THE POST-1984 NETWORK

Same network, new names. The network breaks apart into long distance carriers (AT&T, MCI, Sprint) and local phone companies – the “Baby Bells.” The biggest technological innovation is fiber optics for long distance. Dial up modems offer limited access to the early Internet.

02

EARLY CONVERGENCE / DIVERGENCE

More communications networks begin offering more choices. With cable modems and DSL lines, respectively, cable TV and telecom firms offer broadband access to a common, public Internet, which yields early convergence of data-based services and content. Mobile begins rapid growth.

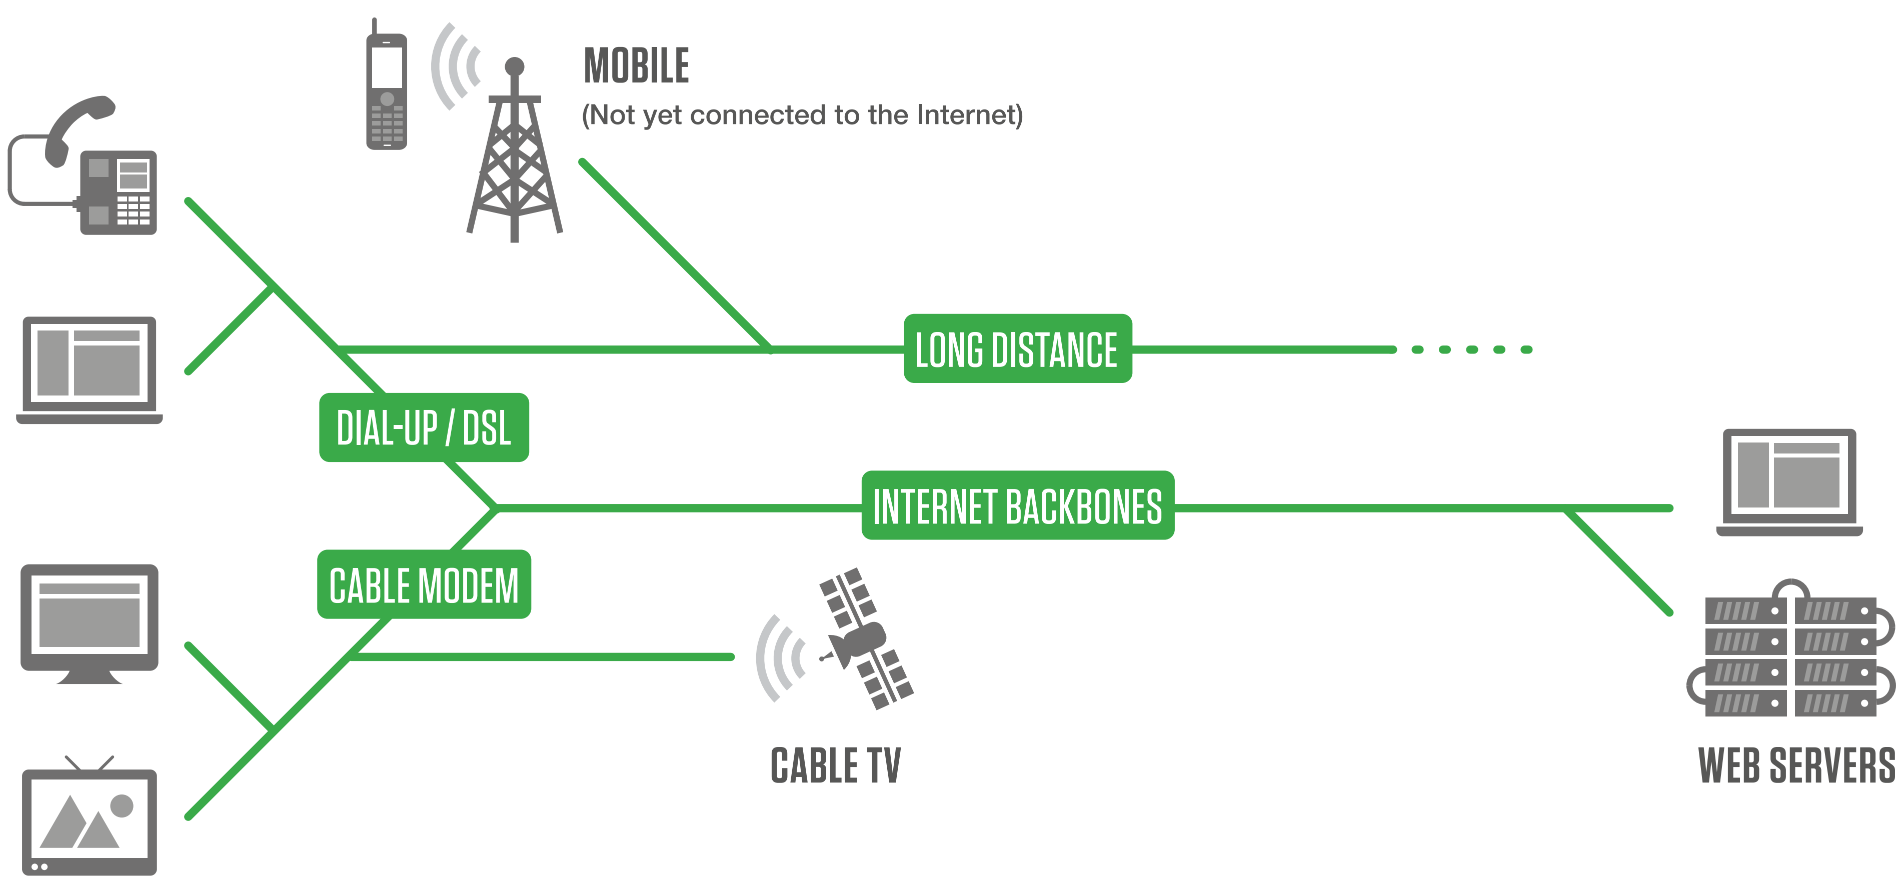

TODAY'S HYPERCONNECTED NETWORK

Five modes of broadband access are supplied by five types of communications service providers. Software, hardware, content, and retail companies become major Internet infrastructure providers, with massive networks and cloud computing capacity, often disrupting older services and media. Mobile devices and the App Economy achieve hypergrowth.

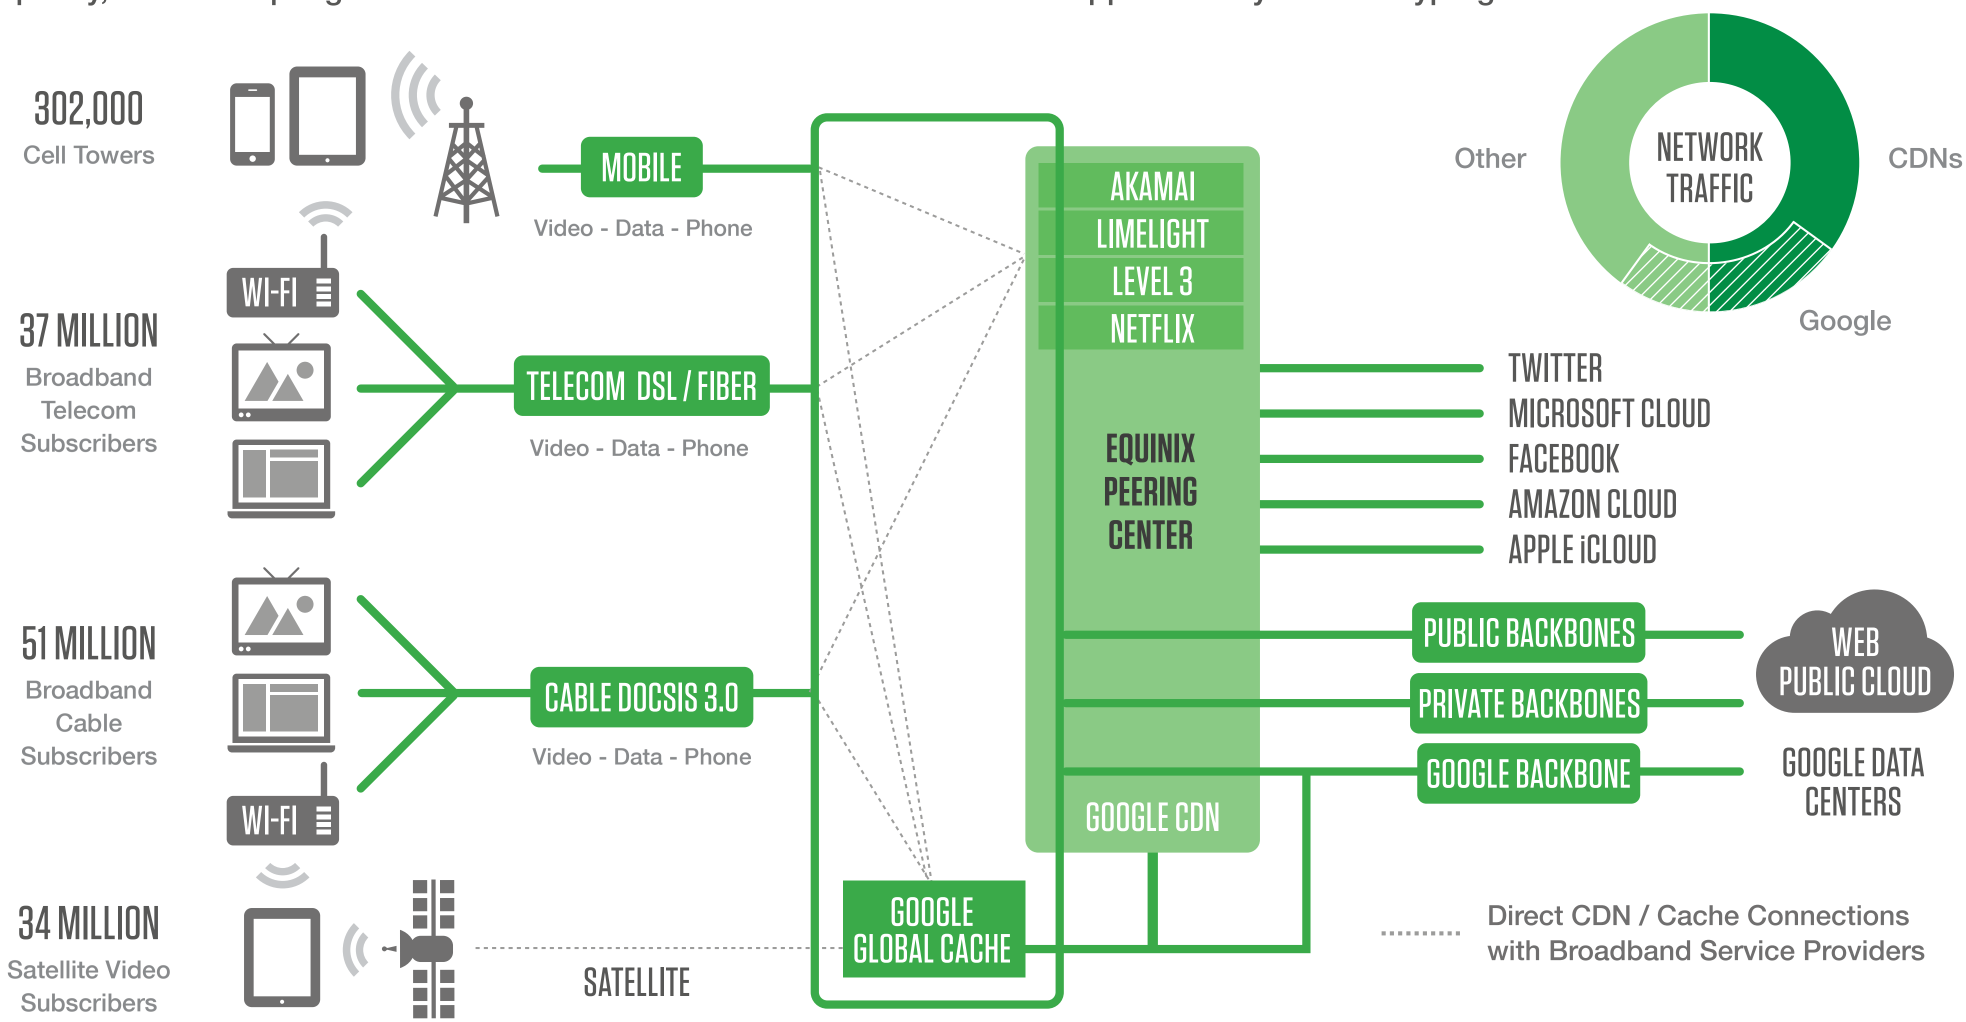

How the Net Works:
A Brief History of Internet Interconnection

BRET SWANSON > February 21, 2014

The Internet is an historic technological, social, and commercial success. It is also a success of self-organization and self-governance. Building something so complex requires exquisite planning by individuals and teams creating the hardware and software to power such a sprawling system. It also requires a conceptual framework that provides just enough commonality to make the pieces work together, but not so much top-down instruction that the system cannot adapt, grow, evolve, and innovate.[1]

We celebrate the Internet's dynamism – most apparent in the ever expanding choices of content, services, and devices that attach to it. Less heralded, but no less important, however, are the networks that power the whole system and the increasingly complex and creative ways all our networks connect to one another.

As the Internet grows in complexity and commercial importance, new network players, new network economics, and new interconnection practices can cause friction among the participants. Some argue we need new laws or regulations to govern the Internet from on high. But with all the industry's positive momentum, abandoning self-governance and commercial give-and-take would be a mistake. The market has proven it will adapt as circumstances change.

We have not reached the end of the line in network innovation. Cloud computing, mobile, real-time telepresence, and other network intensive services will require more bandwidth, more coverage, more connectivity, more up-time, and lower latency, all functions that will require more hyper-connected network capacity. The existing organic process, where engineers and businesses make pragmatic technical and financial decisions, is, in this dynamic environment, far more likely than government mandates to drive growth and accommodate unpredictable innovations.

The Early Internet

A brief history of the Internet helps make the point. In 1969, engineers working on a Department of Defense contract connected the campus computer networks of UCLA, UC Santa Barbara, Stanford Research Institute, and the University of Utah. Arpanet, the seed of the Internet, was born.

Interconnection circa 1998

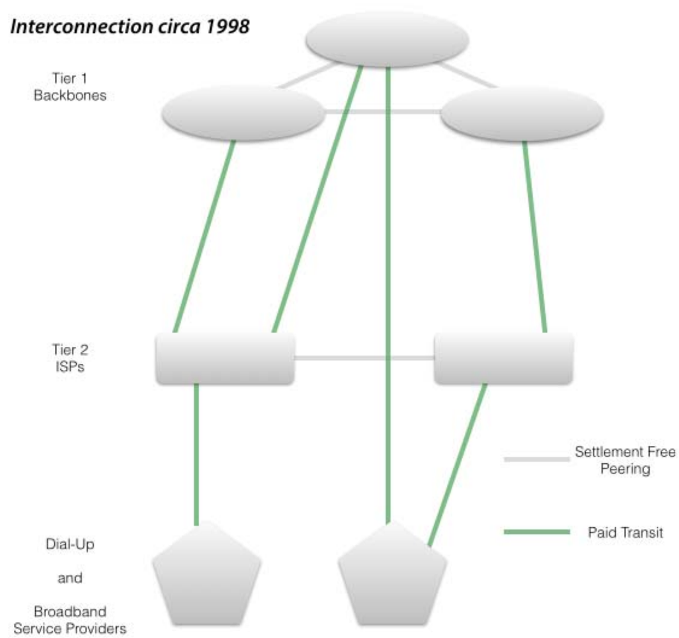

Through the 1970s, more universities and government researchers joined Arpanet, and distinct teams built other experimental networks. Engineers created some of our well known languages and protocols, such as TCP/IP and Ethernet, but they also tried others that did not survive. In the 1980s, the National Science Foundation helped upgrade the backbone network from its original 50 kilobit-per-second telephone lines to faster 1.44-megabit T1 lines, and later to 45-megabit T3s. Private entities, such as UUNET and PSInet, however, also began building backbone networks. We started calling these data networks, collectively, "the Internet."

Getting all these systems to work together was a highly collaborative process. The Internet's early "stakeholders" circulated some one thousand Request for Comment (RFC) memos on protocols and interconnection schemes. In 1984, the domain naming system (.com, .edu, .gov) went into effect, and soon after practitioners from across the globe created two key groups – the Internet Society and the Internet Engineering Task Force – that would help develop the standards and customs that drove the next wave of growth. Between 1985 and 1987 the number of Internet hosts jumped from 2,000 to 30,000, then to 160,000 in 1989, and to one million by 1993.

By the early 1990s, the World Wide Web and Netscape browser shifted the Internet into an even higher gear. In 1990, NSF had lifted commercial restrictions on the NSFNET, and in 1995, NSF privatized it.

Connecting the First Networks

During this period of expanding usage and new, private networks, a number of "exchange points," or network meeting places, emerged. MAE-East, Commercial Internet eXchange (CIX), NSF's Network Access Points (NAPs), and, later, MAE-West and Palo Alto Internet Exchange (PAIX) connected the various networks to one another. These were physical locations where the cables of the various networks connected to allow data traffic to flow from one to another.

This was an unregulated arena, so unlike the world of telecom at the time, with its government-set tariffs, geographic boundaries, and access charges, the Internet players were making up the technical and commercial rules as they went along.

At the exchange points, some of the larger networks with roughly equal traffic flows agreed to trade data traffic at no cost. They called it "settlement free peering," and the choice of words was appropriate. "Peers" were networks similar in size and capability. Because most of the traffic was email, text, and Web pages, traffic tended to be roughly

Interconnection Terms

Tier One ISP — a large continental or global network that, through its own infrastructure and its peering relationships with other networks, can reach any point on the Internet. It does not pay others for transit.

Tier Two ISP — a network, often regional in nature, that connects broadband service providers, content providers and websites, and enterprises to larger Tier One networks. These entities pay Tier Two networks for transit to the Tier One networks, and Tier Two networks pay Tier One networks for transit to the rest of the Internet.

Content Delivery Network (CDN) — a network of computers and "caches" that stores data, webpages, and videos close to end users and optimizes routes across the Internet, both logically and geographically. Content providers and websites pay CDNs to speed their content to end users. Some large content providers like Google have their own CDNs.

Transit — a network access service in which, most often, a smaller entity or network pays a larger entity or network for access to the larger network. Consumers pay their broadband service provider for "transit" to the Internet. Broadband service providers, Tier Two ISPs, and CDNs pay Tier One ISPs for "transit" to the Internet.

Settlement Free Peering — an interconnection agreement in which two networks trade traffic with one another at no cost.

Paid Peering — an interconnection agreement in which networks trade traffic with one another but, because the traffic is "asymmetric" (one network is carrying far more data than the other, incurring higher costs), the party carrying less traffic pays the other a fee to make up the disparity.

symmetrical. Each network was likely to give and receive similar amounts of traffic to the other networks with whom it "peered." Why engage in extra financial transactions with one another if the payments would just cancel out?

Smaller networks and the early Internet access providers like Compuserve and AOL purchased "transit" connections to the larger Internet backbones. These "Tier 2" Internet service providers thus paid to gain access directly to a "Tier 1" Internet backbone and, because the large backbones peered with one another, all points across the Internet. Transit providers could thus be thought of as "ISPs for ISPs."

The First Web Boom

The Internet exploded in the mid- to late-1990s, and its architecture continued to change. Between 1994 and 1996, Internet traffic grew 100-fold, or 10-fold two years in a row. And commercial Tier 1 backbones struggled to keep up. The exchange points were no longer up to the task of establishing enough connectivity, in the right places, in a timely manner. So the backbone networks started to connect to one another in a wider number of large markets using metro area circuits.

Peering politics was sometimes fierce. Networks fought with each other over who was Tier 1 versus Tier 2 and bickered over interconnection terms. Each network carrier wanted, as much as possible, the other networks to connect with it at its preferred location on its preferred terms. (In many ways, this is happening again today.) And yet the market successfully adjusted to the changing environment.

By 2000, a new model was emerging — the large, carrier-neutral, data exchange center. A company called Equinix proposed this new model. It would build large, modern, secure data centers and allow all comers to connect inside its facilities on their own terms. Because it supplied only the meeting space, Equinix marketed itself as a neutral party, a sort-of open super hub for all types of network and content firms. It was a place where you knew all the other networks would have a presence and where, as peering expert Bill Norton described, "large-scale peering interconnections could be established within 24 hours rather than 24 months." [2]

At about the same time, in the late-1990s, two other significant dynamics were changing the interconnection market — broadband access networks for consumers and content delivery networks.

Broadband Access Providers

The cable TV firms grew up serving their customers video content, first via antennas on tops of hills and then via large satellite collectors at their "head-end" facilities in each town or market. The cable firms did not have connections to cross-country or global telecom networks. But the advent of the cable modem meant cable needed a path to the Internet. In the late-1990s, cable's chief links to the Internet were through paid transit arrangements from Tier 2 ISPs such as @Home and Roadrunner.

During the technology crash of 2000, however, @Home failed, and the cable firms began buying transit directly from the Tier 1 backbone providers. The cable firms noticed something else. Much of their traffic was being sent to and from other cable providers. Instead of employing a Tier 2 ISP to reach the Tier 1 backbone, who would then connect to yet another Tier 2 ISP, and then down to the cable firm, why not just establish direct connections with other cable firms?

The broadband service providers — the cable firms and telecom DSL networks — thus began directly exchanging traffic with one another, often inside the new neutral exchange point data centers. Because they were carrying so much traffic within their own customer bases, the larger cable companies, such as Comcast, also began building larger nationwide backbones of their own.

Content Delivery Networks

As the visual Web grew in the late-1990s, content firms, including big dot-coms, news sites, and ecommerce providers, needed to get closer to end users. If an Internet user in New York clicked on a webpage hosted on a server in San Francisco, the content of that webpage would have to traverse the country, often taking indirect routes through as many as 17 router and switching “hops.” (A hop is a physical node on the network — a router or a switch — that data packets touch on the way from origin to destination. More hops mean a less direct transmission, more electronic processing of packets, and ultimately slower and less reliable delivery of packets.) The physical distance and high hop-counts delayed the delivery of packets to the end user and eroded the experience, especially for photos, artwork, banner ads, and other multimedia content. Content providers, who purchased transit through Tier 2 and even Tier 1 ISPs, were dissatisfied.

Akamai, one of the first content delivery networks (CDNs), offered a solution. Replicate and store the most popular webpages and other content in multiple servers, strategically placed geographically and with more closely-coupled connections to broadband access networks. This would reduce both the light speed delay and the hop delay and might even reduce a content provider’s transit bill.

Content firms and websites paid CDNs to get their content closer to end users. CDNs, which consist of tens of thousands of geographically dispersed servers running specialized software that optimizes routes across the Internet, would often pay for multiple high-throughput connections to the broadband providers at strategic points around the country, and around the world.

Few of the early Internet pioneers could have imagined these creative network innovations happening within their conceptual framework, but there were even bigger changes on the way.

Web Video and the Hyper Giants

Launched in 1998, Google, by 2003-04, was growing so fast that it was rapidly taking over entire data centers where it rented space. In 2006, Google acquired YouTube, and with broadband access networks now delivering multi-megabit speeds, Web video exploded. Google needed not just its own data centers but its own content delivery networks and global fiber network. It built them all.

Soon, Microsoft, Facebook, Amazon, Apple, and other content and software firms would do the same. The largest content firms (later dubbed “Hyper Giants” by network scientist Craig Labovitz) had suddenly become some of the world’s largest network firms. This was a silent revolution.

Netflix, the DVD-by-mail company, meanwhile, launched its Web streaming service, and seemingly overnight became one of the biggest bandwidth users on the planet.

The rise of Web video did something else. It substantially altered the mix of downstream and upstream data traffic. Video is thousands of more times bandwidth-intensive than text or webpages, and for movies, sports, and video clips, it is nearly all downstream. That is, end users consume vastly more traffic than they put back into the network.

Transit payments had always been used by smaller networks or content providers seeking connectivity with more end points (that is, seeking to reach a larger audience). And settlement free peering often made sense between similarly situated networks — for example, between two Tier 1 ISPs. But in the past, the traffic and payment flows were simpler and more hierarchical (see network maps on page 5). In general, end users paid broadband service providers and content providers, who paid Tier 2 ISPs, who paid Tier 1 ISPs.

In the new world of YouTube, Netflix, and CDNs, however, an even larger share of the traffic is one-way, at least on many portions

Interconnection, Then and Now — *These figures show simplified network maps, circa 1998 and 2014. Notice the big changes over a mere decade and a half — more players, new connection types, the rise of the "hyper giants," and greater overall complexity. Also notice that the Internet is composed of a mix of paid transit, paid peering, and settlement free peering relationships, among others. (Lines connecting specific firms do not necessarily represent actual network or business relationships. Rather, they show typical connections and business transactions between firms of the type shown — i.e., broadband service provider, Tier 1 backbone, CDN, content firm, etc.)*

c. 1998

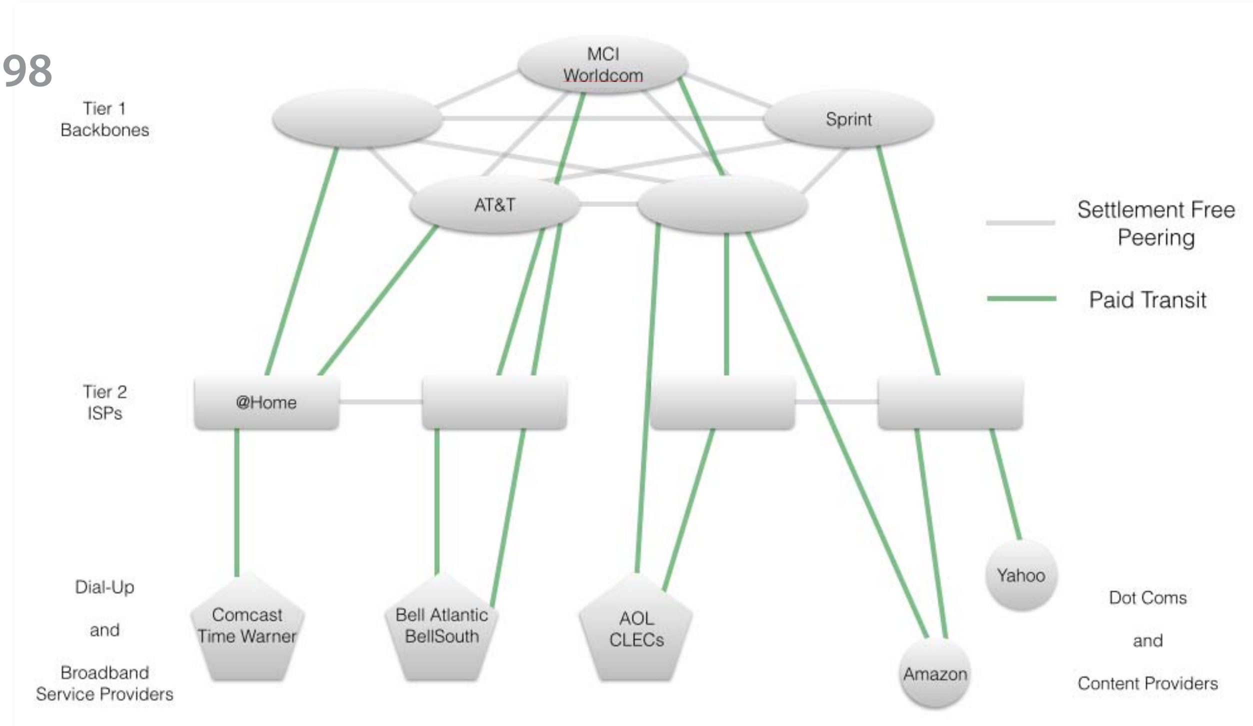

c. 2014

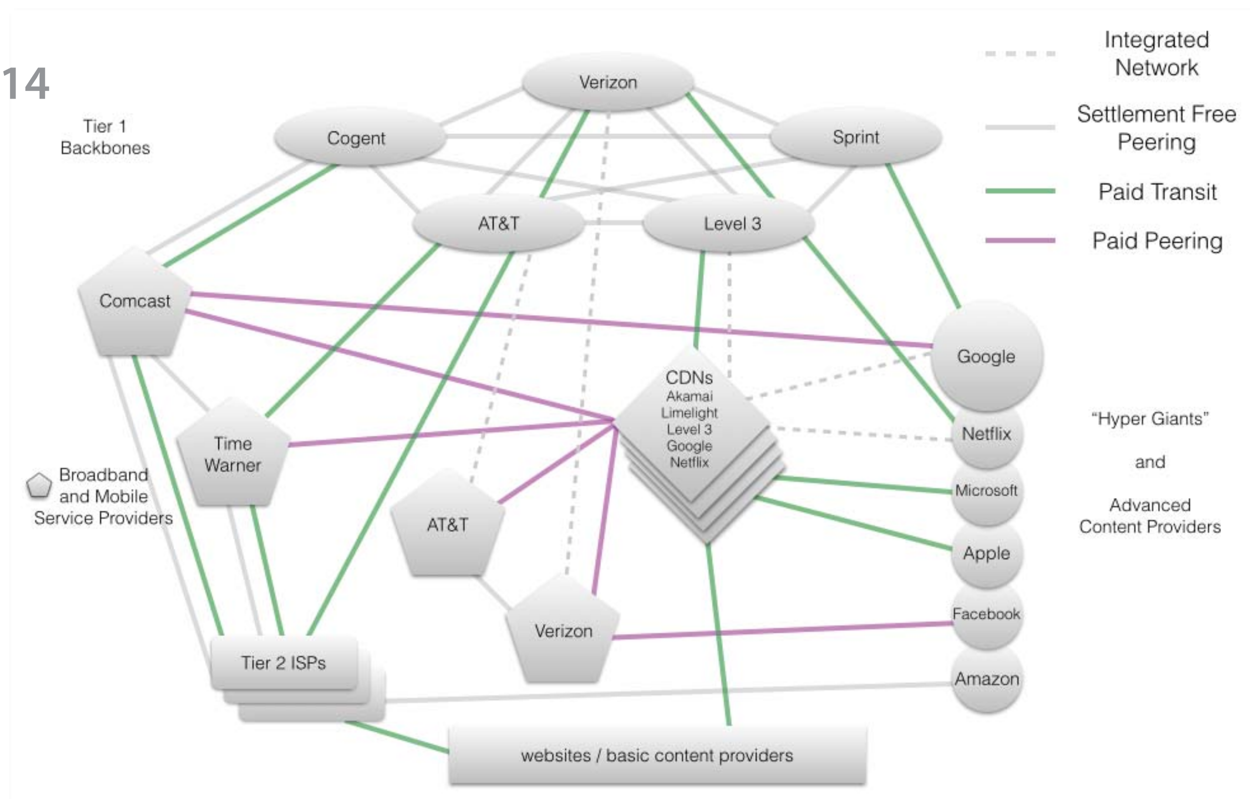

of the network, at many times of the day. And the traffic does not necessarily simply flow "vertically" up to Tier 1 backbone networks and back down. More networks and content providers often connect to one another more directly — or "horizontally" — and in more places (again, see page 5). More networks and content providers thus use more varied and more sophisticated paid transit arrangements and even "paid peering" to account for these highly "asymmetric" traffic flows.

A Rare Public Battle

In 2010, Comcast, Level 3, and Netflix engaged in a high profile battle over the ways Netflix's traffic would reach customers on Comcast's network. Level 3 and Comcast had both transit and peering relationships. And Netflix, through CDNs, had paid Comcast for access. But Netflix and Level 3 had an idea. If Netflix housed its content within Level 3, it could deliver its video to Comcast for free as if it were a peer. Level 3 would enter the CDN business and host the Netflix content for a lower price than other CDNs were charging Netflix to connect to Comcast. Level 3 would get a little extra revenue, and Netflix would cut costs by by routing this traffic over Level 3's settlement free peering links. Comcast would get the downside. Firms reorganize their network operations and business relationships often, and there is nothing wrong with seeking more efficient architectures.

Comcast, however, noticed a significant spike in traffic coming from Level 3 (due to Netflix) and pointed out that this violated its peering agreement with Level 3. Settlement-free peering, remember, had long been limited to situations where networks exchange roughly similar amounts of traffic. Comcast believed Level 3 and Netflix were trying to game the system by exploiting the Comcast-Level 3 peering relationship to dump costs onto the Comcast network. (A network or content firm that mostly sends traffic to others, but does not carry much traffic in return, can impose large financial and network quality costs and upset the economics of the network value chain.) Comcast thus sought to adjust its agreement with Level 3 to reflect this traffic asymmetry. Level 3 and Netflix cried foul, using publicity and regulatory pressure to improve their negotiating leverage. In the end, however, the companies settled on a new agreement, the details of which were confidential — without regulatory intervention.

Considering the number of firms, the complexity of networks, and the pace of change, these episodes have been remarkably rare. The industry is highly competitive but, like most environments free from too much regulation, also highly cooperative.

Ever Changing Interconnection

None of the interconnection arrangements has totally displaced the others. Settlement free peering, Tier 1 and Tier 2 transit, paid peering, and CDNs, among other arrangements, exist side by side. Network relationships and commercial arrangements change according to the quickly advancing technological and financial realities of one of the world's fastest moving industries.

Broadband service providers now even house within their own networks Google Global Cache (GGC) servers, which contain its most highly trafficked content. Netflix, likewise, within the last 18 months, moved most of its video content from third party CDN providers to its own OpenConnect CDN infrastructure. Netflix is also attempting to forge relationships with broadband providers where, like GGC, it would house its content directly within the broadband networks, close to end users.

By 2010, Google's network had grown so large that, according to network scientist Craig Labovitz, it accounted for 6-7% of all Internet traffic. But by 2013, that number paled: Google, says Labovitz, now accounts for up to 25% of the Internet. Netflix, meanwhile, accounts for up to a third of the data flowing over U.S. broadband access networks in evening hours.

Despite the rapid change, tumult, and occasional friction, most of the interconnection world "just works." For example, according to a Packet Clearing House survey of the world's 5,000 ISPs, 99.51% of peering relationships in 2011 occurred without contract, or merely on a "handshake" agreement.

The industry over many decades developed these customs because networks, by their very nature, are highly interdependent. A network that does not have good connectivity to other networks plunges in value. Connectivity is king. The incentives motivate each network player to seek the best service for its customers. ISPs and broadband service providers want their customers to be able to reach as much content as possible, as reliably as possible.

Because of the dramatic changes in content, traffic flows, and the number and type of new network players (the Hyper Giants, for example), the types and terms of interconnection agreements have continued to evolve. Paid transit, paid peering, and other network arrangements will proliferate as the Internet evolves.

The Future

Networks will continue to grow, and interconnections will continue to grow in number and complexity.

Real-time multimedia streams for cloud-based gaming, desktops, and apps will replace many kinds of localized content. These data streams (such as ultra high definition 4K video) will need geographic proximity and, in some cases, interoperability of Quality of Service (or Quality of Experience) regimes that can prioritize content across multiple networks. The delivery of cloud-based apps, services, and content to mobile devices will especially benefit from closely coupled, low-latency links between data centers and mobile access points. (Because a mobile device relies so heavily on the cloud for its computer power and data storage needs — think Siri voice search, Google Docs, or cloud gaming — and because wireless is trickier and more capacity-constrained than is wired, optimizing the links between mobile devices, wireless nodes, and cloud resources can make a big difference in the user's experience.)

Software defined networks will also make new demands on and change the nature of interconnection. Moving network functionality like security, access control, QoS/QoE, remote peering, and network configuration to the cloud will yield large efficiencies and cost savings. Some firms are even considering the centralization and thus virtualization of individual wireless base station functions in remote cloud centers. But these cloud advances will also require big capacity, low latency, and high reliability, straining network performance.

Although asymmetric traffic flows dominated the last decade of Internet content, applications like high-resolution video chatting and conferencing may finally become widespread enough to reverse at least part of that trend, producing more symmetric content.

Whatever the case, all these technologies, products, traffic flows, and business relationships are difficult to predict. The numbers and types of networks will continue to grow, as will the interconnection relationships and overall complexity. Flexibility in network architecture and business relationships is thus crucial to accommodate these innovations.

Conclusion

The Internet is an ever expanding network of networks, where the whole and its constituent parts are ever changing. Where Arpanet linked four entities, each composed of a few end points (primitive computer terminals), today's Internet links thousands of large networks, millions of smaller networks, and billions of increasingly diverse end points (PCs, smartphones, web servers, cloud clusters, cars, and machines and sensors of all types).

To link billions of end points to one another, however, requires organization, cooperation, and trillions of dollars in infrastructure investment. It requires universal standards, like the Internet Protocol (IP), so all the parts work together. But it also requires enough flexibility – in technology, architecture, and commercial relationships – to allow for innovation in networks, content, and services.

From the beginning, our networks have never stopped changing. Nor have the ways networks connect to one another, or the terms. Interconnection disputes are not new, but they have been and remain rare. The size of the Internet economy dictates there will be more disputes (as in any industry), but the industry has and will continue to resolve these disputes in a dynamic, rapidly changing environment, without regulatory involvement. EE

---

[1] The author acknowledges and thanks Verizon for supporting the research in this report.

[2] Bill Norton's website drpeering.net and his books, including *The Internet Peering Playbook, 2013 Edition*, are excellent resources for both the lay reader and the industry insider.